WINTER/SPRING 1994

COMPUTER TECHNOLOGY REVIEW®

THE TECHNOLOGY PUBLICATION FOR THE INTEGRATION/RESELLER MARKET

VIDEO PLAYBACK ON THE DESKTOP

INTEGRATING OCR

SYNCHRONOUS DATA COMPRESSION

FIBRE CHANNEL— WHO NEEDS IT?

VIRTUAL TAPE ON THE LAN

SUPPLEMENT TO COMPUTER TECHNOLOGY REVIEW

To help bring your records storage out of the dark ages and up to speed, get into fast, economical IBM Optical Library Dataservers.

These optical storage systems put thousands of pages of printed reports or microfiche on an optical disk about the size of a CD, then pack as

IBM 3995 Optical Library Dataserver

many as 144 disks into a simple, reliable "jukebox." The result: an online library that locates records in

It sets new records for retrieving old records.

seconds, for scores of simultaneous users—with vastly improved security, too.

IBM 3995 Optical Library Dataservers come in a variety of sizes and configurations. They fit neatly into existing LAN, midsize or mainframe operations, so current applications need little or no modification. And IBM optical drives are designed for optimum reliability.

IBM and IBM Business Partners can provide customized storage solutions, including applications, installation, training and long-term service, too. For more information call **1 800 IBM-6676, ext. 707.**

And put your old microfiche in the storage facility that suits it best: the circular file.

IBM 3995 Optical Library Dataservers store and retrieve millions of pages on write-once or rewritable optical disks, for storage that's as flexible as it is affordable.

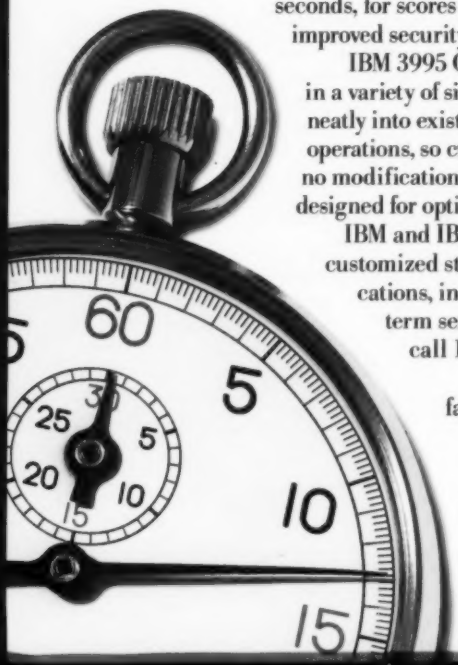

IBM is a registered trademark of International Business Machines Corporation.
©1993 IBM Corp.

IBM

Circle 55 on Reader Inquiry Card

3½" HARD DISK DRIVES

HITACHI'S HIGH CAPACITY 1.4 and 1.0 Gbyte DRIVES

Immediate Delivery

The DK315 has everything you want in a hard disk drive:

- 3½", 1.4 or 1.0 Gbyte capacity
- Fast SCSI-2 interface
- Data transfer rate 4.6 MB/sec. or 2.7 MB/sec.
- 11.8 ms average seek time
- 256-KB read look-ahead cache
- Five year warranty
- 400,000 hr. MTBF

In fact, with its exceptional reliability, the DK315 might just have the lowest cost of ownership around.

Prove it to yourself.
Call 1-800-HITACHI today.

HITACHI®

Five Year Warranty

Circle 43 on Reader Inquiry Card

Immediate Delivery on DK326 1 Gigabyte 1-inch Low Profile

CLARiiON welcomes you

CLARiiON is a trademark of Data General Corporation © 1994 Data General.

to the new Peace Movement.

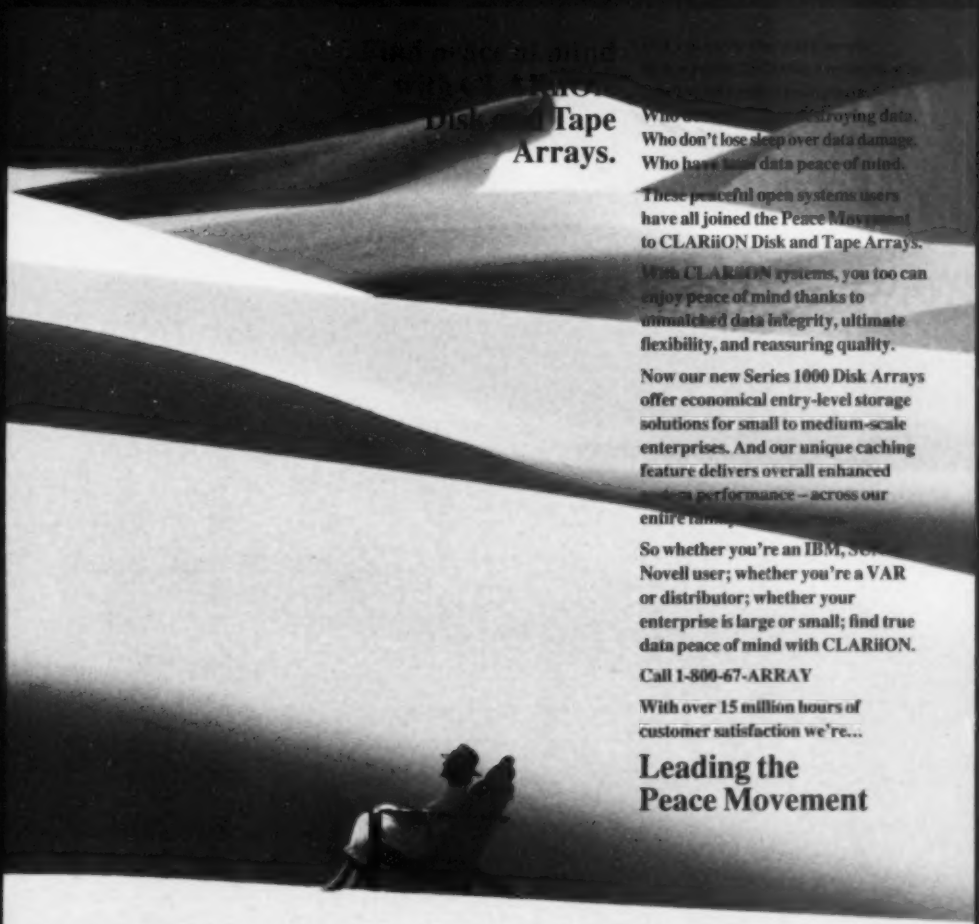

Tape Arrays.

Who don't lose sleep over data damage.
Who have total data peace of mind.

These peaceful open systems users have all joined the Peace Movement to CLARiiON Disk and Tape Arrays.

With CLARiiON systems, you too can enjoy peace of mind thanks to unmatched data integrity, ultimate flexibility, and reassuring quality.

Now our new Series 1000 Disk Arrays offer economical entry-level storage solutions for small to medium-scale enterprises. And our unique caching feature delivers overall enhanced data performance – across our entire family.

So whether you're an IBM, Sun, Novell user; whether you're a VAR or distributor; whether your enterprise is large or small; find true data peace of mind with CLARiiON.

Call 1-800-67-ARRAY

With over 15 million hours of customer satisfaction we're...

Leading the Peace Movement

CLARiiON Advanced Storage Solutions

A Business Unit of Data General The Open Systems Experts

Circle 3 on Reader Inquiry Card

Call Today For Free
Multi-Media Demo Disk
1-800-4-*TeamPro*

A NEW GENERATION OF PC TECHNOLOGY HAS EMERGED!

Introducing *TeamPro*,™ The World's First Multi-PC Workgroup Computer.

TeamPro from HTR is the next metamorphosis in PC technology that matures today's desktop PC into a full-grown workgroup computer. Progressing from inefficient PC-based LANs to a highly-structured workgroup-based LAN; from a slow 10Mbps network transfer rate to a high-speed 80Mbps; from a single CPU to multiple CPUs in one chassis; from single-host devices to multiple-host device sharing; ***TeamPro***'s workgroup technologies prevail. ***TeamPro*** redesigns workgroup connectivity, reshapes network architecture, and maximizes workgroup and network productivity.

TeamPro Is An Expandable PC.

The ***TeamPro*** workgroup computer is a better PC to begin with. Start with one PC and expand by adding CPU cards as the workgroup grows. A tower-sized ***TeamPro-4000*** can accommodate up to 5 different ***TeamCPUcards*** from 386/486 to Pentium, providing tremendous cost savings and flexibility in system integration, expansion and upgrades. In addition, any existing PC can be converted into a 2-PC workgroup computer by using the ***TeamPro-II*** add-in kit. ***TeamPro***'s unprecedented built-in workgroup capability has made it the new desktop standard for the 90's.

TeamPro Is A Workgroup Computer.

Coupled with a peer OS, like NetWare Lite™ or Windows for Workgroups™, the ***TeamPro*** workgroup computer is the total system solution for a workgroup covering an area of up to 20,000 sq. ft. Featuring specially designed workgroup-oriented hardware and software, ***TeamPro*** provides such benefits as:

- instant plug-and-play workgroup connectivity at 80 Mbps
- multiple-host device sharing that optimizes overall performance
- self-diagnostic capability for multiple CPUs
- cost savings from eliminating redundant hardware
- "green computing," with reduced noise and low power consumption.

TeamPro offers superior workgroup computing that cannot be achieved by loosely-connected PCs or a small peer-to-peer network.

TeamPro plug-and-play workgroup computing.

TeamPro Is A Network Subnet.

As the PC-based LAN evolves toward the Workgroup-based LAN, the ***TeamPro*** workgroup computer becomes the perfect network building block. Used as a network subnet, ***TeamPro*** provides distributed processing for the workgroup while reducing unnecessary network traffic to the file server. ***TeamPro*** improves network productivity and security by:

- instantly creating multiple network nodes
- allowing concurrent workgroup communications
- preventing possible data loss due to heavy data transfer
- isolating outdated equipment from slowing down the network
- excluding unauthorized access to the file server.

The ***TeamPro***-based LAN makes workgroup computing and client-server computing work together perfectly.

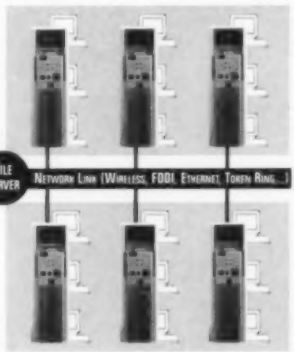

The better-architectured ***TeamPro***-based LAN.

TeamPro Is Available Today.

The future of networking is available today in one cost-effective, powerful solution. HTR invites you to explore business opportunities with us, and see how we can expand your product offering. Our VAR programs offer a healthy profit margin and extensive technical support. HTR is strongly committed to the success of each of our VAR partners. Contact us today.

1-800-4-***TeamPro***

HTR™

HT Research Inc., 1342 Bell Avenue, Tustin, CA 92680
(714) 566-9100, FAX (714) 566-9109

Circle 4 on Reader Inquiry Card

HTR and TeamPro are trademarks. HTR's product concepts and workgroup-oriented components are currently patent pending. All product names are trademarks of their respective companies. Copyright © 1993 HT Research, Inc.

COMPUTER TECHNOLOGY REVIEW®
THE TECHNOLOGY PUBLICATION FOR THE INTEGRATION/RESELLER MARKET

Publisher
YURI R. SPIRO

Associate Publisher
CAROL L. STAGG

Editorial Director
GEORGE McNAMARA

Senior Editor
MARK C. FERELLI

Associate Editor
DAVE TROWBRIDGE
trowbridge@genie.geis.com

Editorial Manager
SUSAN GRISANTI

Production Manager
EDMOND WONG

Production Coordinator
JOHN HAYASHI

Technical Illustrator
CLAUDIA GRAY

Circulation Director
MARIA WONG WHITE

Circulation Assistants
CARLOS BARRAZA
DUANE KING

Advertiser Services
ABBY NELSON
Inquiry Handling

SALES OFFICES

National Sales Manager
STEVEN L. TANEMAN
310/208-1335 ext 26

Pacific Northwest
LYNN S. GIORGI
310/208-1335 ext 19

Midwest/Southwest/So. Atlantic
CAROL L. STAGG
310/208-1335 ext 16

West/Southwest
CLARK EMERSON
310/208-1335 ext 38

Mid-Atlantic/New England,FL
ROBERT A. POGGI
516/288-8380

Mountain States
STEVEN L. TANEMAN
310/208-1335 ext 26

Card Deck
East Coast
CARLOS GUTIERREZ
310/208-1335 ext 27

West Coast
ABBY NELSON
310/208-1335 ext 36

COMPUTER TECHNOLOGY REVIEW
THE TECHNOLOGY PUBLICATION FOR THE INTEGRATION/RESELLER MARKET
CORPORATE AND SALES HEADQUARTERS
924 Westwood Blvd., Suite 650
Los Angeles, CA 90024
(310) 208-1335
Printed in U.S.A.

Contents

8 THE BUSINESS OF TECHNOLOGY

10 IN FOCUS

13 DESKTOP INTEGRATION

21 NETWORK TECHNOLOGIES

41 CD-ROM / IMAGING / MULTIMEDIA

51 OPTICAL STORAGE

59 DISK STORAGE

91 TAPE STORAGE

COMPUTER TECHNOLOGY REVIEW®

THE TECHNOLOGY PUBLICATION FOR THE INTEGRATION/RESELLER MARKET

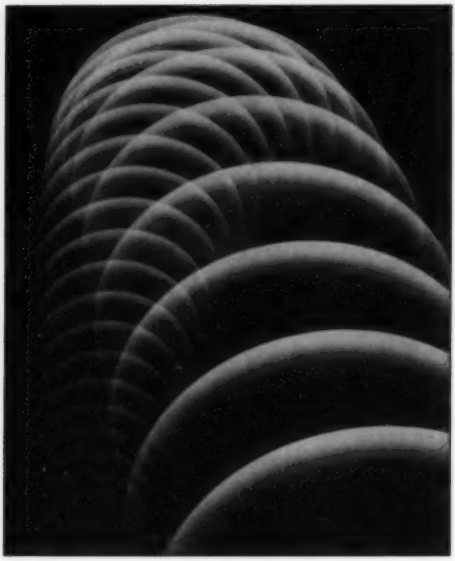

The only thing constant is change.

Change—technological change—is what drives the market and creates the conditions for the continuing growth of the computer industry.

COMPUTER TECHNOLOGY REVIEW has covered changing computer technology since our inception in 1981. During this period we have been changing, too. In this issue you will see new evidences of this in fresh design and improved readability.

But what won't change is our dedication to the enduring value of technology. It's that dedication that keeps CTR the only technology publication serving the integration/reseller market.

DESKTOP INTEGRATION

14 OCR

OCR ALONE CAN'T ADDRESS ALL DOCUMENT PROCESSING NEEDS

Adding new "intelligent" technologies to OCR could eliminate data entry as we know it.

RANDALL STERN
Symbus Technology Inc.

18 CROSSPOINT ARCHITECTURE: NEXT STEP IN MULTIPROCESSING

AARON BOXER
Concurrent Computer Corporation

NETWORK TECHNOLOGIES

23 SYNCHRONOUS DATA COMPRESSION EMERGES AS WAN COST-CUTTER

SDC over dial-up lines can duplicate switched 56 (SW56) or ISDN functionality for as little as one-third the cost.

DAVID McNAMARA
Motorola Codex

29 NETWORK UPS

FERRORESONANT UPSs FACE PFC PROBLEMS

AARON L. DAVIS
American Power Conversion

32 NETWORK SECURITY

A LAYERED DEFENSE PLUGS THE VIRUS SECURITY GAP

SUSAN RICHARDSON and RONALD SULLIVAN
Reflex Inc.

THE BUSINESS OF TECHNOLOGY

"A company, if not careful, can re-direct itself in both marketing and engineering to the point of winning the current price battle, but losing the war on a company-wide basis."

LEE ELIZER

IN FOCUS

"Unlike too many networks today, networks should be easy-to-use and inexpensive to maintain and easy to expand."

DENNIS HAYES

34 NETWORK STORAGE

SEPARATING FACT FROM FICTION IN SMS (STORAGE MANAGEMENT SERVICES)
SMS is a set of specifications designed to solve the problems of managing data in heterogeneous environments.

MICHAEL KIDD
Palindrome Corp.

CD-ROM / IMAGING / MULTIMEDIA

42 HIGH RPMs ADD SPEED TO CD-ROMs AND MINICHANGERS

RICHARD MILLETT
Pioneer Technologies

46 CD-ROM

MPEG VIDEO CD STANDARD HASTENS DIGITAL VIDEO PLAYBACK ON THE PC
The pathway for adoption of digital video on the PC will become very clear in 6-12 months.

HOWARD GORDON
Xing Technology

OPTICAL STORAGE

52 STANDARDS

THE EASY ROAD TO INFORMATION INTERCHANGE OPTICAL STANDARDS
The growing list of optical standards takes some of the mystery out of the technology.

JEFFERY GARNER and ARNOLD JONES
Micro Design International

DISK STORAGE

60 FLASH MEMORY

CHOOSING AND INTEGRATING CURRENT FLASH SOLID STATE STORAGE SYSTEMS

MARKUS LEVY and PEDRO VARGAS
Intel Corp.

66 RAID

STREAMING RAID—A DISK ARRAY MANAGEMENT SYSTEM FOR VIDEO FILES
Conventional file servers and storage systems are not well-suited to support video over LANs.

FOUAD TOBAGI, JOSEPH PANG, RANDALL BAIRD and MARK GANG
Starlight Networks

84 PCMCIA

BLOCK FLASH FILE SYSTEMS DELIVER MASS STORAGE FOR MOBILE COMPUTING
By emulating a disk at the sector level, Block FFS makes flash memory an ideal solution for mobile mass storage.

AMIR FRIEDMAN
M-Systems Flash Disk Pioneers

TAPE STORAGE

94 LAN BACKUP

VIRTUAL TAPE BRINGS COOPERATIVE APPROACH TO LAN BACKUP
Virtual tape functions as a CPI-C/LU6.2 gateway between LAN backup software and the mainframe.

GREG COTICCHIA
Emprise Technologies

105 FIBRE CHANNEL

DEFINING FIBRE CHANNEL FOR NETWORK AND CHANNEL PROTOCOLS
Fibre Channel is a new I/O interface for network and channel users to expand interconnects for tape, disk and other peripherals.

JAN DEDEK and GARY STEVENS
Ancot Corp.

Lee Elizer
of Data Storage Concepts

THE BUSINESS OF *on* TECHNOLOGY

Price War No Reason To Reduce Research And Development

Alternate technology battles are a subtle form of a price war. Exabyte, the leading and only manufacturer of 8 mm tape drives, has been experiencing a price war from the alternate technology of 4 mm DAT for the past few years. The net effect of this alternate technology price war (albeit several other features such as a 3.5-in. form factor were also contributors) has been a deterioration in 8 mm prices and a leveling to decline in revenue and shipment growth of 8 mm drives versus 4 mm. But, important to the end user, Exabyte has increased its R&D spending by more than 50% both as a percentage to revenue and in absolute dollars during the past year.

Already the 4 mm DAT drive market has been reduced—partially due to price wars—to two main suppliers (HP and Conner), with three smaller suppliers—Sony, WangDAT and Exabyte. This is from a total of 36 companies developing DAT drives just four years ago.

DAT drive product prices continue to drop, almost on a quarterly basis reminiscent of disk price-war policies. The effect on 4 mm DAT R&D has been to accelerate developments for higher performance products, while at least one company is gambling on off-shore manufacturing for cost competitiveness.

The high performance .25-in. full-size cartridge drive market has Wangtek and Tandberg sharing the majority of product shipments. These two companies are relatively financial equals, albeit not in the same league as HP or Conner that have dropped this type of product in favor of alternate technologies.

The primary price war for the full-size .25-in. cartridge will come from alternate technologies: 4 mm DAT, 8 mm and/or DEC's DLT drives. This lack of a direct price war is particularly enhanced with the recent agreement between Wangtek and Tandberg for standardization on future high-capacity versions of full-size .25-in. cartridge drives.

R&D Financing Dynamics

The initial reaction to a price war is to assume that as a company's profits decline the company will not have (or apply) enough resources to continue to be competitive on an R&D basis.

Quite the opposite is often true. A company will over-invest in R&D when profits decline in the belief that it can engineer its way to profitability. This to the detriment of an effective market presence. The obvious mistake is to try to engineer to profitability in the same product line that is involved in a price war. There is a point of diminishing returns where engineering could best be spent in the follow-on product rather than cost-reducing the current, price-pressured product.

An alternative R&D disease is a company controlled by "financial" types who arbitrarily keep R&D spending at 5% or less in order to pump up short term profits. This strategy is successful for a relatively short period of time, but then a relatively long period of time may be required to return to parity with the competition.

Normally, R&D expenses are treated as a direct expense on the profit and loss statement for a tape drive manufacturer. Normal R&D spending is in the 6-8% of revenue range when a company is at an acceptable profit margin. But will a company increase R&D direct expenses when its product growth begins to slow or decline?

Another approach to R&D spending is balance sheet financing. The advantage in this approach is to attempt, depending upon accounting principles, to keep the R&D expense from appearing on the profit and loss statement, thereby raising short-term profits. An example of this approach is a financial investment in a joint venture, where the joint venture performs the R&D for a product in return for future considerations. This balance sheet approach is protected from the income statement in some instances, allowing for continued, or increased, product development without negative P&L results.

Off-balance sheet financing is a third approach to R&D spending. This approach involves a company owning a portion of a joint venture without a significant financial investment. Another company—perhaps an R&D partnership—develops the product, the other company markets the product and perhaps manufactures the newly developed product in return for consideration. This was a popular approach in the early 1980s. StorageTek, as an example, was involved in a number of R&D product partnerships during this period.

The Technology Impact

Naive companies cease to plan and strategize during a price war. Alternatively, a company, if not careful, can re-direct itself in both marketing and engineering to the point of winning the current price battle, but losing the war on a company-wide basis.

Historically, if a company has diversified its product lines, it will survive and prosper. Those companies that focus strictly on market share, and have a limited product portfolio, experience wide swings in profitability as competitors continue to drive prices down.

R&D continues to be expended during price wars, however, only not necessarily in the existing product line. The technology advantage to the end user during a price war thereby expands, as more new products are introduced.

Price wars are healthy... at least for the end user. ■

Lee H. Elizer *is president of Data Storage Concepts, Boulder, Colorado.*

YES, YOU C'AN TAKE IT WITH YOU.

If you're going mobile, you should be taking Seagate along. Introducing the Seagate FlashCard—the ultimate storage solution for the new age of data portability.

Whether it's for a palmtop, notebook, or even a desktop PC, here's a tiny device with some huge advantages. It's extremely rugged. It's easy on your batteries. And when you take it out and take it along, the FlashCard allows you to pack-up your office and put it in your pocket.

Seagate's FlashCard complies with the universally accepted standard, PCMCIA-ATA. And you can choose from a range of models from 1.8 to 20.9 Mbytes of formatted capacity.

PC Card Call 408-438-8111 if you'd like more information—or if you're ready to go.

Seagate

THE DATA TECHNOLOGY COMPANY

Circle 5 on Reader Inquiry Card

Seagate Technology's Flash Memory Mass Storage is designed and manufactured by SunDisk Corporation. SunDisk and the SunDisk logo are trademarks of SunDisk Corporation. Seagate, the Seagate logo and the Storage by Seagate logo are registered trademarks of Seagate Technology, Inc. PCMCIA is an acronym for Personal Computer Memory Card International Association. © 1993 Seagate Technology, Inc.

In Focus

WITH DENNIS HAYES
Hayes Microcomputer Products Inc.

Your network should work hard to let you work easy.

The computer industry should shift its thinking to "Is the LAN as we know it today what we need to be offering our customers?" Are networks what they should be in order to serve our customers and to take the customers applications into the future?

Unlike too many networks today, networks should be easy-to-use, easy and inexpensive to maintain and easy to expand. To obtain these goals, you need a power-ful network that is designed to be more complicated on the inside so that it is simpler on the outside. This type of networking technology cannot be just another layer of technology, another version, laid on top of existing technology, a process made so popular in recent years. We must look at a new wave of technology that will take us to the next level of connectivity.

Today's Network Perspective

As an industry we tend to view LANs from a product-oriented, technology perspective. We talk about high-end networks and low-end networks. High-end LANs are considered to be powerful, have advanced features, and everyone knows they require constant maintenance and an expensive network administrator who constantly functions as a network mechanic. We usually look at the client-server model when referring to high-end networks.

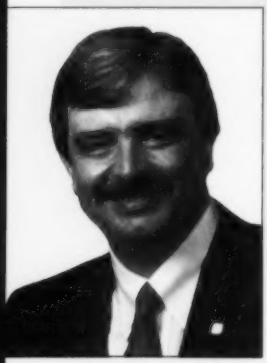

When we talk about low-end LANs, we think of the DOS-based, peer-to-peer networks that are generally easy to install but lack many of the advanced features of high-end networks. Today's low-end LANs still require a large degree of ongoing maintenance particularly once they begin to grow in size above a single server or a few workstations.

The User Perspective

It is increasingly important for the industry to look at networks not from a "hardware/software" perspective but from the users' perspective. New users who have a very limited technical expertise but who want the benefit of the network are constantly entering the market. These large numbers of market-entry users just want the benefit of the network and are not really interested in the technology.

Customers only know what they want: to be able to easily connect with other people to share infor-mation—usually locally, often remotely. And, they do not want to hire a network mechanic.

The traditional approach has been that if you fit that customer profile, you start out with a DOS-based, peer-to-peer LAN and sacrifice capability. At some point as your needs grow, you go through the conversion from a low-end network to a high-end network and pick up the cost of maintenance for that high-end network and hire a network administrator.

What Networks Should Be

Installing a network should be fast and simple. It should be "self-configuring," using a series of simple questions to determine the best settings for a particular user environment, giving the installer flexibility without excessive overhead.

For a network to be easy to use, you should not have to use complicated commands or file names. You should be able to quickly access information and resources on the network. You should not have to know where a file is stored. That is the LAN's job. You told the LAN to put something on the server, the LAN should remember where it put that item and be able to retrieve it for you when you need it. The same is true for shared printers. You should not have to use complicated commands to access them. The network should do that for you.

Networks should offer an easy, simple way to expand the number of users on the network. It should not inter-fere with the normal operation of the network or impact other users on the network when you add users. You should not have to take down the network, re-install or reconfigure to implement network upgrades.

People should be able to demand a lot from their network; the network should not demand a lot from its user or administrator.

Distributed Server Networking

Distributed Server Networks such as that used in Hayes LANstep overcome the difficulties associated with Central Server or Peer-to-Peer networks, including the time it takes to manage these networks. This technology gives you the flexibility to distribute the network workload over different machines, allowing the fine tuning and load balancing to fit your unique environment, but without sacrificing the ability to use the resources of the network easily, regardless of where they are physically located.

The end result is that the entire LAN is linked as a single computing environment that gives LAN users the flexibility to move from distributed management to central management, or some combination of both. And the benefit to this arrangement is scalability. When you triple the size of your network everything doesn't get three times slower and nine times harder to manage.

Distributed networking is a much more powerful environment than the earlier approaches and at the same time it is possible for us to provide automated and streamlined management capabilities that make LAN management easier.

In short, distributed Server Networks are what today's networks should be. ■

Dennis C. Hayes *is president and founder of Hayes Microcomputer Products Inc. (Norcross, GA).*

Sun Microsystems now come with these extra moving parts.

They're Merisel's Sun Team and you'll find these dedicated experts running all over the place.

To help you.

Because for configuring, supporting, maintaining or expanding Sun-based solutions into new markets, you won't find a better team of highly trained technical professionals under the Sun.

We have Systems Engineers who will work closely with Sun VARS, offering invaluable UNIX expertise on all Sun Microsystems workstations. From entry-level Classic Workstations and servers to the powerful 20-processor SPARC Center 2000. You'll also find that Merisel has the broadest line of complementary UNIX hardware and software anywhere.

Our Sun Team will help you with a dedicated 800 number to instantly access all the UNIX sales and tech support you need. Backed up by monthly newsletters, technical bulletins, training schedules, catalogs, pricing on diskette and customized price books.

They'll offer specialized financial services. Including net terms, flooring, hardware and software leasing, standby letters of credit, and cross corporate guarantees that'll make your money work even harder. And if you think that's great, you won't believe all the extra service and support you'll enjoy by joining our exclusive VAR Vantage program.

So call us today at 1-800-SUNDIAL and take advantage of the clout and expertise that only the Full-line Channel Specialist can offer.

And let our moving parts get your business moving.

MERISEL
World Class Distribution

Call for details about Sun products on our On-Line Literature Library at (310) 415-9886. All product names are registered trademarks of their respective corporations. ©1994 Merisel, Inc. Corporate Offices: 200 Continental Blvd., El Segundo, CA 90245 - (310) 615-3080

Circle 65 on Reader Inquiry Card

MAXION™ Multiprocessor System... The MAXIMUM Real-Time Solution

Real-time computing is tough, really tough. Real-time situations are typically "do or die," "must have," "time is money" situations. Your information system must be totally reliable, rugged, and provide you with real-time access to unlimited data about real-world events. Concurrent is the undisputed innovator in Real Time.

Concurrent's astounding new MAXION multiprocessor system and UltraSMART™ Scalable Multiprocessor Architecture for Real-Time are maximized to overcome performance bottlenecks. The MAXION multiprocessor system incorporates the revolutionary "crosspoint" processor interconnect architecture that eliminates standard bus system limitations. Another Concurrent innovation.

Super-fast at 150 MHz, and compact, with a VME-6u double-height European form factor, MAXION multiprocessor systems are ideal in demanding environments. Simulation and training, weather detection and airspace management, measurement and control, signal intelligence, financial trading, and healthcare — areas where Concurrent excels.

To learn more about the innovative MAXION multiprocessor system and UltraSMART, call Concurrent at **1-800-631-2154** or write to Concurrent Computer Corporation, Dept MAX, 2 Crescent Place, Oceanport, NJ 07757.

MAXION and UltraSMART are Trademarks of Concurrent Computer Corporation.
© 1991 Concurrent Computer Corporation.

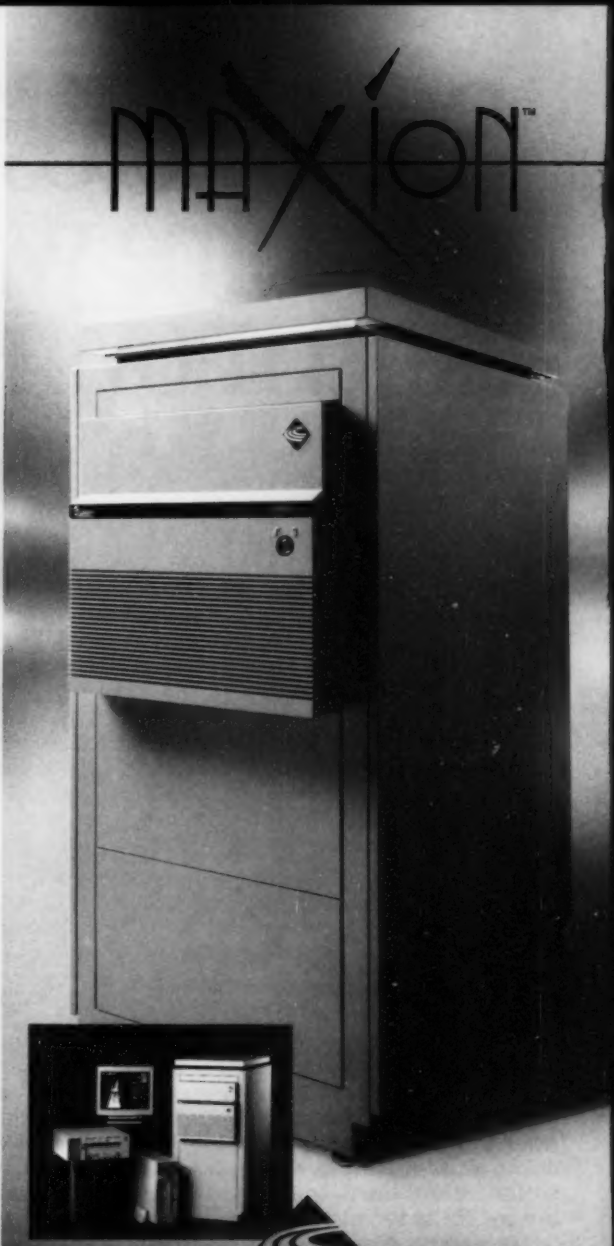

CONCURRENT COMPUTER CORPORATION™

THE REAL-TIME PEOPLE.

Circle 6 on Reader Inquiry Card

DESKTOP INTEGRATION

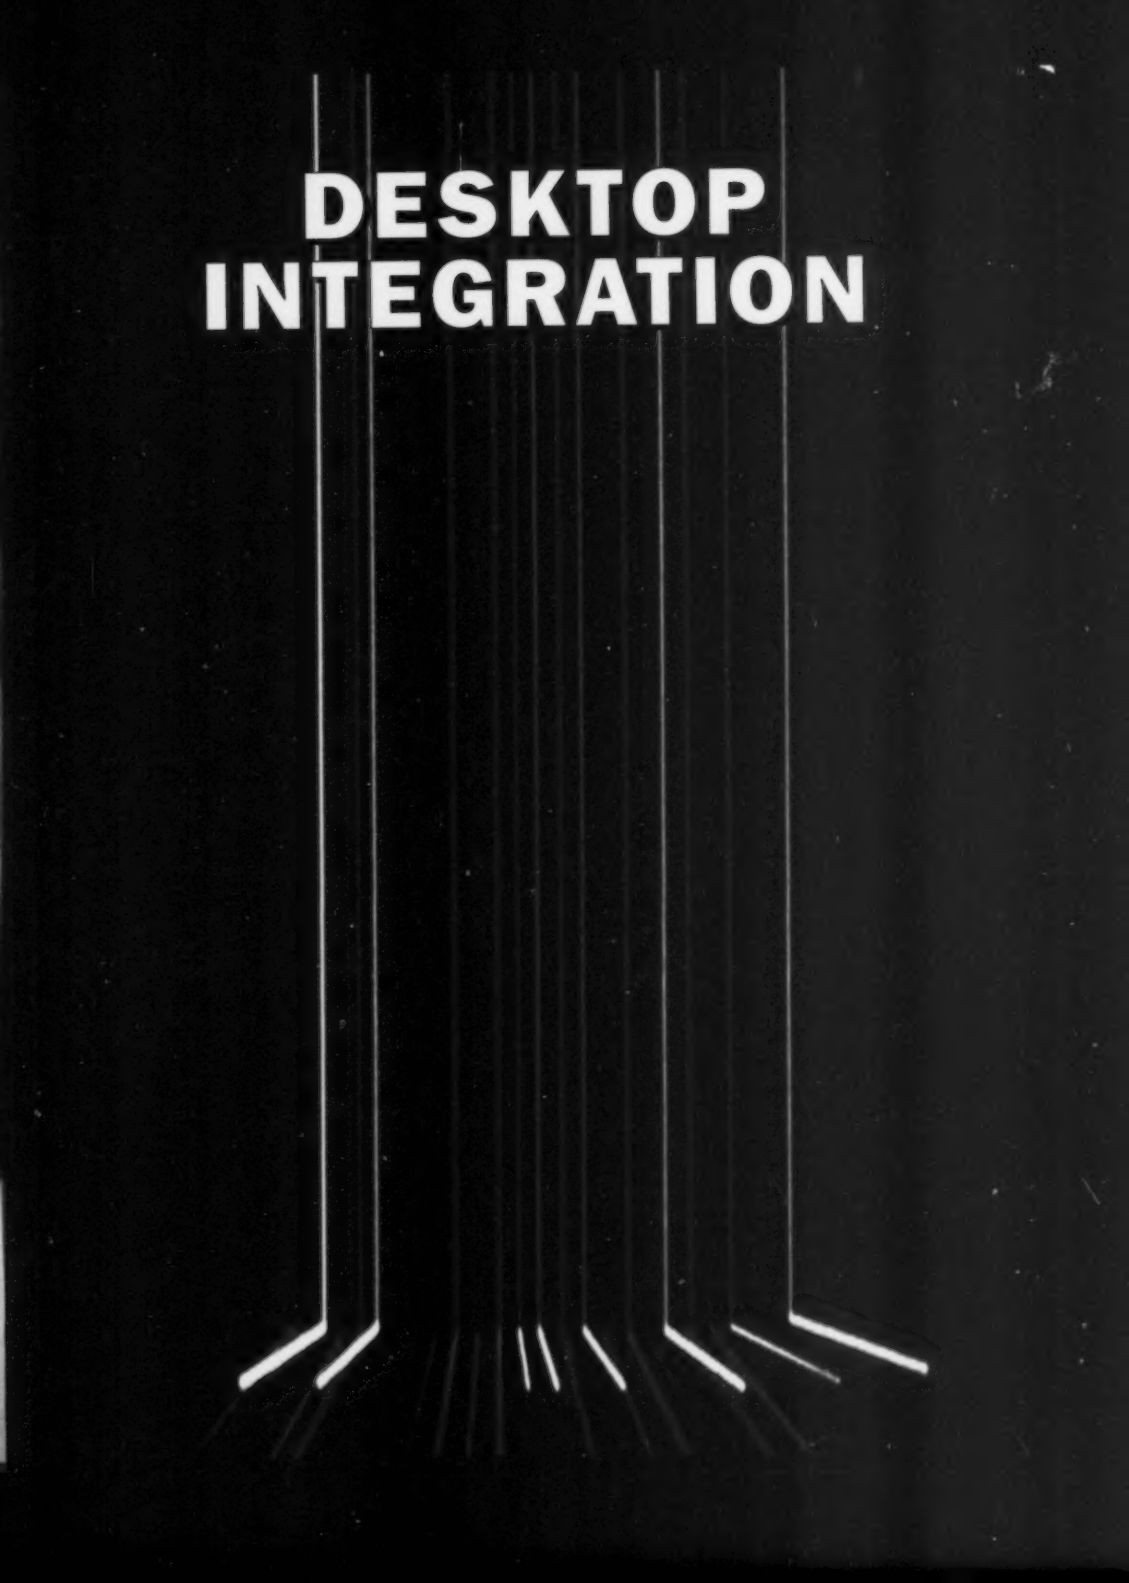

OCR

OCR Alone Can't Address All Document Processing Needs

Adding new "intelligent" technologies to OCR could eliminate data entry as we know it

Optical Character Recognition (OCR) is one of several technologies that many people associate with document image processing. Under this universe, we define "imaging" as the capture, display and storage of images in digital form.

OCR systems are one tool for selecting and delivering portions of digital images. OCR recognizes machine-printed text and converts it into ASCII files. It works optimally with large amounts of machine-printed text, such as in scanning technical manuals or entering legal cases into a litigation support system.

OCR systems cannot read hand-print entries, or mark-sense ("check-box") entries. They often fail at recognizing degraded ma-chine print, as in fax copies, or poor quality photocopy reproductions. Applications where documents have frequent ink smudges, dirt or skewed images also will confuse an OCR system and cause errors. By themselves, OCR systems are inadequate for forms processing.

When an image is unclear, an OCR system will make its best determination and input the value. It will flag low-confidence characters and words that do not appear in a lexicon, but the system has no way of detecting incorrect results—errors that are far more costly to correct. The higher the volumes, the more that OCR data integrity becomes an issue.

For example, a system that can correctly identify nine out of 10 digits of a long distance telephone number field is helpful, but still yields a wrong number. If the telephone number was cross-referenced with a lookup table of area codes, an incorrect character in the area code would be detected, and the entire phone number made correct.

OCR vendors readily admit that OCR products will make errors, but position OCR systems as far more effective than retyping. While this is true, many applications cannot tolerate the random input of erroneous data. Issuing checks for the wrong amount, for example, is clearly unacceptable.

Some newer, low-volume products boost data integrity by adding word pattern modules or morphology tools. Word pattern modules identify known patterns such as month/ date/ year entries, social security numbers, or phone numbers that may appear in any location on the page. Morphology software dissects a printed page to determine headlines, captions, wrap-around copy —and better enable the OCR engine to present text in the correct page reading order.

These tools can improve results in reading general text, but are hampered by a lack of knowledge of exactly what these patterns are expected to look like. They also do not address the relative sensibility of a value—something that is essential in automating data entry because it is the element that a data entry clerk would provide.

Another issue is OCR's tendency to confuse the form image with the data. To compensate, many OCR users print forms using special "drop out" inks that scanners cannot read. These add an extra printing

by Randall Stern
Symbus Technology Inc.

200 IPS × 8.75" = DMT2820

Fastest Quad-Density GCR in a Slim Profile

Anritsu lets you record data at speeds of up to 200 ips with the new DMT2820 GCR Streaming Tape Drive. That gives you host data transfer rates of up to 2 MB/s via SCSI asynchronous interface or 1.4 MB/s via Pertec interface.

Yet this high-performance tape drive is extremely compact, just 8.75 inches high. It will save you space in any rack configuration.

The DMT2820 half-inch tape drive lets you choose between quad-density recording speeds of 800, 1600, 3200, or 6250 cpi, and uses the NRZI, PE and GCR recording methods for maximum flexibility.

A 2 MB buffer memory provides better control of data transfer rates and data block size (up to 64 kbytes) with automatic error recovery. The residual tape length is clearly displayed so you can easily keep track of the residual memory capacity.

High-speed tape rewinding at 320 ips completes a full 2400-foot tape in just 90 seconds or less.

The DMT2820 also features automatic loading and threading, a set-up menu, self-diagnostics and other easy operating functions. The front panel display shows the current internal settings, and lets you know instantly of any error conditions. All maintenance can be performed from the front of the unit.

Lower speed requirements can also be satisfied with the **DMT2720**, which records at 100 ips.

Whether you're operating a minicomputer or a mainframe, the Anritsu DMT2820 offers a fast solution to memory back-up that won't crowd your available space.

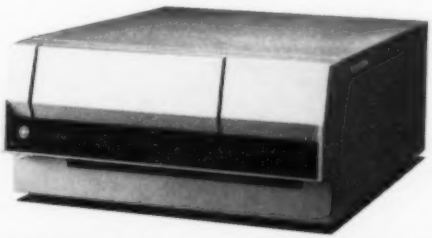

DMT2820 GCR STREAMING TAPE Anritsu

ANRITSU CORPORATION 5-10-27, Minamiazabu, Minato-ku, Tokyo 106, Japan Phone: 81-3-3446-1111 Fax: 81-3-3442-0235
ANRITSU AMERICA, INC. 365 West Passaic Street, Rochelle Park, NJ 07662-3014 Phone: 201-843-2690 Fax: 201-843-2665

Circle 63 on Reader Inquiry Card

OCR

TECHNOLOGY	CAPABILITIES	BENEFITS	WEAKNESSES
OCR	MACHINE-PRINT TEXT RECOGNITION DOCUMENT RECOGNITION	● AUTOMATES FULL-TEXT ENTRY ● INDEXING BY CONTENT ● PRESERVES COMPLEX LAYOUTS	● NO HAND PRINT ● UNSURE IF WRONG ● DEGRADED PRINT IS A PROBLEM
IMAGE PROCESSING	CAPTURE AND STORAGE	● SAVES ON STORAGE ● SEND OR ROUTE FAST ● CAPTURES TEXT, GRAPHICS	● MANUAL INDEXING ● NO EDIT, RE-LAYOUT
ICR	HAND-PRINT TEXT DEGRADED MACHINE-PRINT	● HELPS MORE APPLICATIONS ● BETTER CHARACTER ACCURACY	● NEEDS MORE FOR FORMS PROCESSING ● NO SCRIPT
FORMS PROCESSING	ALL THE ABOVE PLUS FORM REMOVAL, VALIDATION, DATABASE LINKAGE	● NO PRE-SORTING NEEDED ● NO SPECIAL FORM PRINTING	● NO SCRIPT

step and raise costs. More importantly, they require complete control over the creation and provision of forms. Some OCR systems use so-called form template software tools that match the form image against a preset selection, then remove it. However, slight photocopy reductions, the variations caused by multiple printing sites, or the changing of just one field can invalidate the form removal—and decrease processing throughput.

Image processing enhances images for better viewing and printing before recognition by an OCR system. It can clean up fuzzy characters, smooth dot-matrix print, or filter dirt from a noisy background. Image processing also can perform some levels of de-skewing of images.

Most applications can benefit from image processing to filter imperfections that could introduce errors in recognition. Limited image processing functions are incorporated into most OCR systems, and can be purchased separately as well. For example, cleaning-up the smudged background on a telephone book, and performing de-skewing and de-slanting functions all require image processing.

One hardware-based image processing technique is the use of digital signal processing (DSP) chips to accelerate the process of enhancing images. Dramatic advances in DSP chip performance at falling costs have caused some experts to suggest that all motherboards will contain a DSP within the next several years. This will allow every system to perform sophisticated image processing without the need for added accelerator boards.

No matter what its speed, however, image processing does not permit editing or re-layout of documents. It merely facilitates the more efficient storage, viewing and routing. In addition, image processing requires manual indexing, and does not support full-text searches or automated entry into a database.

Intelligent Character Recognition

ICR is based on powerful pattern recognition technologies that can recognize hand print and other patterns with greater variability than OCR—including degraded machine print. Note that "hand print" is different from handwriting; no system exists today to accurately and efficiently recognize unconstrained script handwriting.

ICR systems are considered intelligent because they use a neural network that functions similarly to the human brain. The neural network simply recognizes patterns of signals—in a fashion similar to the neurons, synapses and dendrites of the human brain—and classifies them into pattern classes. It learns to segment and recognize characters more efficiently, with even higher accuracy, the more it is trained.

The neural network will do a "fuzzy" match to recognize unclear or degraded machine print, and compensate for skewed images. ICR systems also incorporate other technologies to find characters, parse them into properly bounded fields and recognize words.

ICR systems are a significant improvement upon OCR capabilities for many applications. They enable the recognition of faxed documents, for example, whose degraded characters are too variable for accurate recognition by OCR alone. They are useful in very uncontrolled applications that have a great deal of variability, but in a defined context. For example, ICR systems will accurately read postal codes and bank remittance printing at the bottom of checks.

Forms Processing

Forms processing tools recognize and validate highly variable characters—both hand- and machine-print—on forms that contain many fields. The need to deliver meaningful data from insurance claims, income taxes, police records, and other complex forms to a central database have made

RESELLERS!

For A Unique
Network and Workstation
SOLUTION

Turn back to our ad On Page 4

1-800-4-*TeamPro*

HR

this a fast-growing market for integrators.

Fig 1 shows the use of OCR, image processing, ICR, and other functions within a typical forms processing system. Forms processing tools work well in applications where an organization cannot control the exact appearance of the form—or the way the user fills out the form. These tools are essential for sustained system throughput when application volumes are high.

For example, the government-standard HCFA-1500 form mandated for use in processing Medicare claims actually exists in over 100 different variations—each posing a unique recognition challenge. That challenge multiplies when health care providers put a coffee cup on the form, write a check mark which overlaps onto an adjacent field for the zip code, and then submit the form for processing.

Forms processing systems improve productivity by drastically reducing the hours that an operator must spend in rekeying hand-print or unclear machine print. Because of the high volumes—over 1,000 or more documents daily—often seen in such applications, forms processing systems incorporate added validation steps to assure throughput and minimize operator intervention.

These steps include:

- form identification—so users can process different forms at the same time, without the need to pre-sort them into batches;
- separating the form image away from the data—eliminating almost any restriction on form layout, conformance to a template, or size;
- image processing at the level of specific fields. Fields, which are made up of properly bounded characters, contain the meaning, and thus the value, of the information contained on a form. Newer forms processing systems view only those fields that the user defines as important. Then they clean up the fields for maximum accuracy and system throughput. Special processes can be applied for hand-print, dot matrix and degrade machine print recognition—all of which can be detected on the fly;
- recognizing hand- or machine-print in designated fields;
- parsing important embedded data at designated fields;
- verifying the validity of the field, not just the string of characters, through the use of context-based rules processing. This allows the wisdom of the human operator to be applied each time a field is recognized. For example, context-based rules will send a field to a clerk for verification if it suggests that a zip code is invalid for the state listed on the form;
- linking the data to a database elsewhere in the organization;
- allowing for minimal human verification when needed. Forms processing systems route images to an operator for manual verification when the system's confidence in the image falls below a preset value. That value can be made very high, in the case of dollar-value fields—or low, on fields where accuracy is important but less paramount.

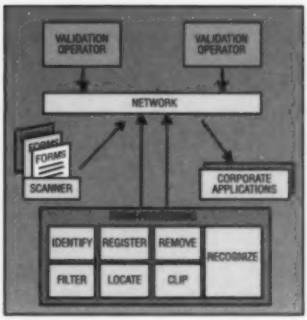

Fig 1 The InScript forms processing system unites OCR, image processing, ICR and forms processing functions.

Instead of requiring a clerk to spend eight hours a day in re-keying data, forms processing systems do the bulk of the data entry—and require only one hour or less from the clerk each day. When orchestrated by a savvy integrator, we have seen forms processing systems deliver labor savings of up to 85%—and paybacks of under a year, based on labor savings alone.

"ICR is based on powerful pattern recognition technologies that can recognize hand print and other patterns with greater variability than OCR."

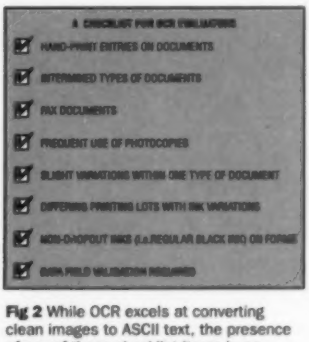
A CHECKLIST FOR OCR EVALUATORS
HAND-PRINT ENTRIES ON DOCUMENTS
INTERMIXED TYPES OF DOCUMENTS
FAX DOCUMENTS
FREQUENT USE OF PHOTOCOPIES
SLIGHT VARIATIONS WITHIN ONE TYPE OF DOCUMENT
DIFFERING PRINTING LOTS WITH INK VARIATIONS
NON-DROPOUT INKS (i.e.REGULAR BLACK INK) ON FORMS
DATA FIELD VALIDATION REQUIRED

Fig 2 While OCR excels at converting clean images to ASCII text, the presence of any of these checklist items in an application will require the use of additional technologies for adequate performance.

Is OCR Enough For You?

Today is an outstanding time for buyers of volume business automation solutions because the choices multiplied several fold. Fig 2 lists the major considerations in determining whether OCR is enough for an application's needs. The challenge is to be as informed as possible about choices, and to know the cost/benefit implications of the various increments of technology. ■

Randall Stern *is vice president of engineering for Symbus Technology Inc. (Waltham, MA). He received a BS in mathematics and computer science from SUNY at Stony Brook, and an MSGG from M.I.T.*

For more info circle 202.

Crosspoint Architecture: Next Step

by Aaron Boxer
Concurrent Computer Corporation

In the last five years, many developers of real-time applications have turned to platforms based on loosely coupled multiprocessors to achieve throughput, response, computation power, and determinism. These systems need close synchronization of independent isolated processes and the low software overhead of communication through shared memory. Designers of general purpose systems are not yet addressing the real-time requirements for synchronization and communication among processing nodes in distributed environments.

The architecture of the MAXION multiprocessor real-time system from Concurrent Computer Corporation addresses the growing mismatch between the rate at which processors demand data and the rate at which memories can supply it. In the last decade, processor instruction speeds have gone up by a factor of more than 80. DRAM speeds are up by a factor of only about two. And the bandwidth of affordable buses has not increased sufficiently for the needs of multiprocessor systems in which processors compete for access to a central memory (Fig 1). Bigger caches to reduce bus traffic has been the traditional solution. But this doesn't work well in real-time applications where frequent interrupts and context switching drive up the miss rates regardless of the cache size.

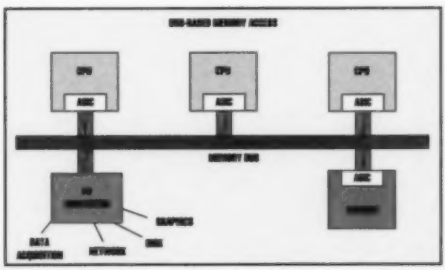

Fig 1 The bandwidth of affordable buses has not increased sufficiently for the needs of multiprocessor systems in which processors compete for access to a central memory.

In the architecture Concurrent chose for its multiprocessor real-time system, there is no single central memory which the processors access via a bus. Instead, system memory is distributed with each CPU, physically as well as logically (See Fig 2). On each CPU/memory module, system DRAM memory is less than six inches away from the CPU. There is only one ASIC of logic between the R4400 RISC processor and DRAM. Memory latencies between a CPU and its local memory are limited by the speed of DRAM, not by the access time of a shared bus.

Applications can be written for this system so that almost all of a process's out-of-cache memory accesses go to the CPU's local memory, and the operating system has a tuning feature that maximizes local memory usage for applications not written to take advantage of local memory. But CPUs sometimes still need to access other CPU/memory modules' memories, and they need to access I/O devices on the VMEbus.

To provide for this, the system's CPU/memory modules are connected by a six-way crosspoint interconnect, implemented with ASICs, rather than a traditional bus. The crosspoint supports up to four CPU/ memory modules and up to two VME64 I/O buses. The crosspoint is a true multipath crosspoint. There is not just one set of wires inside the ASICs connecting all

In Multiprocessing

the ports like a bus. Instead there is a separate 64 bit wide connection between each port, so that four memory reads can be in progress over the crosspoint simultaneously. The miss penalty of the system design equals a bus based design for remote memory accesses and is twice as fast for local memory accesses. The overall bandwidth of the crosspoint is several times greater than the bandwidth of a bus of the same width because four CPUs can be using it simultaneously.

In a bus based multiprocessor system, the memories and buses are rarely idle and contention for these resources increases the cache miss penalty that each CPU sees. Multiple processors generate enough traffic to cause serious contention in a bus based system, to the extent that adding processors beyond two or three brings rapidly diminishing returns. The system's crosspoint isolates much of this traffic, providing linear scaling of performance with the addition of processors, as well as improved determinism.

The requirement to support familiar standards meant that the system's distributed memory architecture couldn't break the global shared memory model that is assumed by many applications and that is familiar to programmers. This requires that the address space seen by each CPU includes all the memory on the other CPU modules. It also requires that cache coherence among all the CPUs is maintained for all the memory. Bus based systems usually maintain cache coherence with bus snooping schemes. Every processor's cache keeps track of which memory locations it is sharing with other caches and for which ones it has the only correct copy of the data. It watches all the read/write traffic on the bus to maintain the correct status of this data in its cache. This takes away cache bandwidth that the CPUs could use. But in the MAXION multiprocessor real-time system, not all of the traffic can be seen by each CPU because of the crosspoint isolation, so a directory based method of maintaining cache coherence is used. In directory based cache coherence there is, logically, a single directory outside the CPUs' caches which keeps track of which memory locations are in each CPU cache and whether one cache has the only correct copy. Physically the directory is implemented as DRAM in the memory system for lowest cost and board space. The system directory cache coherence scheme reduces the coherence traffic, frees up cache bandwidth without adding duplicate tag stores for each cache, and makes the read/write traffic isolation of the crosspoint possible.

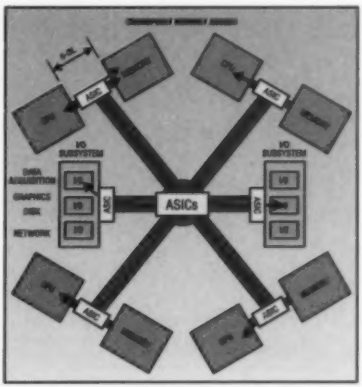

Fig 2 The crosspoint memory access architecture allows multiple simultaneous transfers.

Aaron Boxer *is senior consulting engineer at Concurrent Computer Corporation (Westford, MA).*

For more information circle 210.

2,317 associates, 852 manufacturers, 65,151 resellers, 14,988 products, and someone else is still No. 1.

You.

Our business philosophy is based on putting others first.

Precisely 65,151 at last count.

Which is the number of resellers who know that when we say we're looking out for No. 1, we're talking about them.

It's a practice partly responsible for putting Merisel on the Forbes 500 and Fortune 100 fast-growing companies list. And why so many customers and industry analysts call us the world's best distributor and integrator of computer-based products, services and information.

Here's what we're doing for you.

Products. With 14,988 products available from 852 manufacturers, you know you'll find whatever you need and lots more at Merisel. Macintosh, DOS and UNIX. Printers and peripherals. Systems and software.

Prices. Along with competitive prices, at any given time we're running right around 80 promotional offers. So your cash goes right where it belongs, namely into building your business.

Service. With over 260 hours a year of constant training, your personal Merisel sales rep is specially trained to understand your business needs. Of course, once you've placed your order, you can get it right away from any of our nine U.S. warehouse locations.

Support. We know that the microcomputer products business is getting more complex. That's why we offer some of the most flexible financing options in the business. Tech support with no equal. Great educational opportunities. And innovative on-line services.

Products. Prices. Service. Support. At Merisel, we know what's important.

But we also know that the most important part of our business is your business. To find out how important, call us at 1-800-MERISEL.

Even if you become our 65,152nd reseller, you'll always be No. 1.

MERISEL
World Class Distribution.

© 1994 Merisel, Inc. Corporate Offices: 200 Continental Blvd., El Segundo, CA 90245 - (310) 615-3080

Circle 60 on Reader Inquiry Card

NETWORK TECHNOLOGIES

YOU DEMANDED... WE DELIVERED!

INTELLISERVER

ETHERNET COMMUNICATIONS SERVER

IntelliServer is a compact, high-speed, "plug-and-play" Ethernet communications server that lets you instantly connect up to 64 PCs, workstations, dumb terminals, printers, and modems to a TCP/IP LAN. Start with 16 ports, then snap on 8-port and 16-port expansion modules whenever you need them. There's even an RS-422 expansion module for widely-distributed topologies.

IntelliServer installation and operation are menu-driven. Each user can start up from one to eight sessions and quickly "hot-key" from one session to another. ***Every IntelliServer port is PPP capable*** for inbound/outbound communications, with full-duplex line speeds up to 200K bps!

To ensure maximum throughput, IntelliServer includes a 15 MIPS RISC CPU, 2 MB RAM, and 512K ROM. BNC and AUI connectors are standard, permitting connection to Thin-Net, Thick-Net, or Twisted-Pair media. All popular industry-standard protocols are supported, including *rlogin, telnet, telnetd,* reverse *telnet, rsh, rcp, ping,* PPP, SNMP, SLIP/CSLIP, DNS, and ARP/RARP.

Available in "Tower" and "SlimLine" versions, the IntelliServer is backed by a 5-Year Warranty, free Technical Support, INTERNET access, and a 24-hour BBS with free software upgrades.

Computone also offers the world's largest selection of multi-port adapters and long-distance multi-port cluster controllers. **For free information, call Computone VAR Sales at 800-241-3946 x230 or 404-475-2725 x230, Fax: 404-664-1510.**

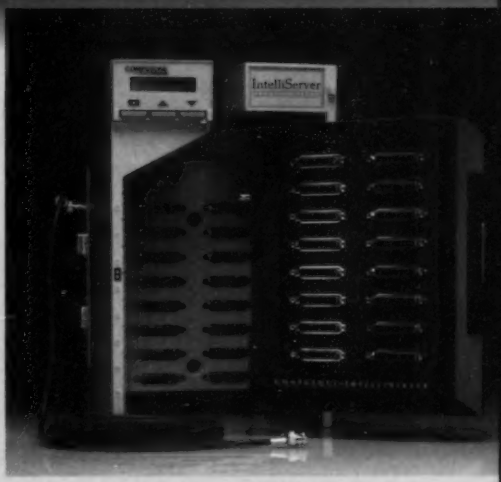

Ask about our Competitive Trade-In Specials!

Circle 7 on Reader Inquiry Card

COMPUTONE
Corporation

TEN YEARS OF INTELLIGENT MULTI-PORT!

Synchronous Data Compression Emerges As WAN Cost-Cutter

SDC over dial-up lines can duplicate switched 56 (SW56) or ISDN functionality for as little as one-third the cost.

Synchronous data compression represents one of the most significant advancements in bandwidth efficiency gains in the past five years. As the demand for high-speed switched and dedicated synchronous connectivity in the LAN internetworking and branch networking markets explodes, SDC technology brings markedly improved performance and significantly lower operating costs.

Combining SDC with V.fast technology, as Motorola Codex has done in its 326XFast-SDC modems, brings the economy and simplicity of the PSTN to users who otherwise would be forced to employ switched 56 (SW56) or ISDN services at one-and-a-half to three times the operating costs (Fig 1). Because the combination of SDC and V.fast modem technology operates over standard dial-up voice lines, savings for Tariff 12 customers can be even more dramatic. Thus, SDC brings the switched versus dedicated operating cost crossover point within reach for new applications like dial-up routing and bandwidth on demand. Such applications can take advantage of the "pay-as-you-go" tariff advantages of voice grade dial up lines.

With SDC, dedicated circuit users can gain in several ways. For example, typical DDS circuits cost between 1.3 and 2+ times more than the alternative analog leased-line circuits that users now can employ using SDC products. Even more significant savings are found with the short distance "feeder" circuits, which can account for as much as 60% of network operating costs. Deployment of SDC products on analog lines in this portion of the network can cut these operating costs by as much as two-thirds. SDC technology, when employed on analog circuits, can provide branch networking access at speeds of up to 72 Kbps, more than sufficient for most of today's feeder networking needs. This gives analog leased-line users a very attractive alternative to migrating their applications to higher speed digital services.

Cost savings of such service upgrades are even more dramatic for users of private microwave links to remote or offshore sites. Instead of consuming 4-5 Frequency Division Multiplex (FDM) slices of a microwave link to provide 56K service, SDC products can simply replace the existing analog modems consuming the same microware link bandwidth while providing throughput up to 72 Kbps. This can result in savings in microwave equipment replacement, reconfiguration and digital equipment deployment.

And what about digital transmission facilities? When SDC is incorporated in a 56K DSU, the bandwidth improvements and network savings can be staggering. For example, Motorola Codex's recently announced 3512 SDC DSU delivers synchronous transmission speeds up to 256 Kbps over a standard 56K DDS leased digital circuit. The only way to provision this kind of bandwidth without SDC is with Fractional T1 (FT1) services, available only from interexchange carriers (IXC). Monthly FT1 services can cost three to four times what a 56K DDS circuit costs and that's on the IXC portion only (Fig 2). Since local exchange carriers don't offer T1

by David McNamara
Motorola Codex

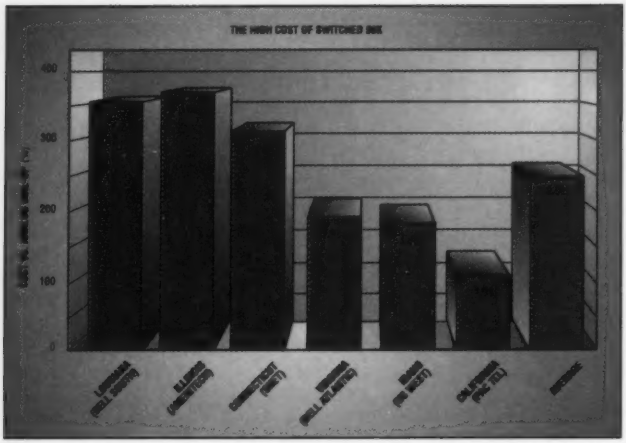

Fig 1 Switched 56 tariffs differ widely throughout the United States, making the ability of SDC to deliver similar data rates over the PSTN quite attractive.

services, the local access portion of a FT1 circuit is provisioned and tariffed as a full T1 line even though the delivered service is only fractional T1. FT1 installation charges and FT1 DSU/ CSUs aren't cheap either.

Network access, per minute usage fees and monthly leased-line charges are obvious "hard" cost advantages of SDC. However, significant hidden or "soft" cost savings also are gained through the simplicity of deployment and complete management facilities that come with equipment employing SDC.

Synchronous Data Compression: What Is It And How Does It Work?

Like asynchronous data compression, SDC must deliver the exact same information content presented at the transmitting station. But there are unique differences between synchronous and asynchronous data compression. For example, in asynchronous data compression, the data must be delivered in sequence. However, idle times between characters can be omitted or added without affecting the integrity of the received message. With SDC, there are additional constraints. In particular, frame integrity must be preserved. That means the HDLC/SDLC frames must be delivered in sequence with all the original information and without introducing any inter-character idle time. It also is important that each frame be compressed and transmitted without introducing significant throughput delay, either in the transmitter/compressor or the receiver/decompressor.

To achieve these criteria, the SDC compressor must recognize the boundaries of the HDLC/SDLC frame, compress only the contents of the frame, transmit it in sequence and deliver it to the receiving entity intact with little or no added through-put delay. Because SDC is of the lossless class of compression and the transmission media can be error-prone, the transmission link between the compressor and

CIRCUIT DISTANCE MILES		1HZ KBPS FRACT. T1		ANNUAL SAVINGS PER CIRCUIT WITH 56KBPS & SDC
50		$1,935		$13,860 TO $16,740
100		$1,980		$14,160 TO $17,160
250		$2,100		$15,000 TO $18,600
500		$2,325		$16,740 TO $21,000

- AT&T ABOS TARIFF AS OF AUGUST 1993
- INCLUDES AVERAGE OF 10 MILE LEC LOCAL ACCESS CHARGES
- DOES NOT INCLUDE INSTALLATION, DISCOUNTS, AND EQUIPMENT COST

Fig 2 SDC over switched 56 can deliver up to 256 Kbps throughput, yielding Fractional T1 throughput at a staggering savings, since the local access portion of a FT1 circuit is provisioned and tariffed as a full T1 line.

decompressor should be error-protected to prevent loss of synchronization. For this reason, the SDC link should be protected by an error-correcting protocol.

As an example, the SDC feature employed in Motorola Codex's 326XFast-SDC modems and 3512-SDC DSU utilizes the CCITT V.42 error-correcting protocol together with a modified version of the CCITT V.42bis data compression algorithm to provide error-free compressed data transmission. These devices further employ an Adaptive Packetizing Algorithm (APA) to fragment the large HDLC/SDLC frames to optimize the throughput on the transmission link while minimizing throughput delay (Fig 3). If the transmission link is experiencing errors due to analog line disturbances, the APA feature fragments the HDLC/SDLC frame into smaller packets to improve the probability of error-free transmission. When the line disturbances subside, APA increases the packet size to minimize packet overhead and thereby maximize throughput. This fragmentation further complicates the pro-cess of guaranteed frame integrity, but improves the throughput of the link while minimizing throughput delay. This insures that even under the harshest line conditions, the HDLC/SDLC frame is delivered in sequence error-free with maximum throughput and minimum through-put delay. To minimize further the throughput delay of an SDC link, the SDC products from Motorola Codex employ continuous compression and packetization in the transmitter and provide several user-selectable options for receiver throughput delay minimization.

But what if the data is not compressible? Suppose it's partially compressible? To deal with this, an expansion avoidance algorithm is employed that automatically monitors whether the data is compressible or not. If the internal compression monitor detects that the data is not compressible, data compression is turned off dynamically to insure that expansion does not occur. When the expansion avoidance algorithm detects the presence of compressible date, it switches the compressor back on to insure that the link is utilized at the maximum

1000 Floppies, 20 SyQuests or 1 OPTIMA DAT.

With DeskTape software, the Optima DAT appears right on your desktop.

Just click and/or drag to open, close, print or transfer files.

Now get the Simplest, Fastest access to 2GB on the desktop for storage-intensive applications.

With Optima's "Just sell and forget about it" DeskTape software and DAT drives for your Macintosh Customers!

- ☑ ANIMATION
- ☑ VIDEO
- ☑ DTP
- ☑ CAD

It's Desktop DAT That's right, **TWO GIGABYTES**, (about a thousand floppies or twenty 105Mb SyQuests™) on a single, fast access cartridge (under 20 seconds average for a full 2Gb tape). And our DAT acts just like a hard drive ... open, close, print, and select and drag objects right into (and out of) it, directly, without restoring...even run QuickTime movies!

Have your customers transfer large files, up to 2Gb, **on a single 125 DAT cartridge**... or recommend they use one to archive the entire year. Better yet, have them add one to their network to free up space on all their drives.

There's no cheaper or better solution (and it doesn't look like there will be). That's why so many end users are switching to DAT right now.

Give us a call today! Ask for a free DeskTape demo, the software that puts any certified DAT drive on your customer's desktop.

OPTIMA™ TECHNOLOGY

Call 714/476-0515 ext. 288 now for more information. (or fax to 714/476-0613).

All trade names are registered trademarks of their respective owners.
*Based on average price of 105Mb SyQuest cartridge.

Circle 42 on Reader Inquiry Card

efficiency possible.

Unsuitable Networking Applications For SDC

SDC technology provides significant bandwidth efficiencies on all compressible traffic that is HDLC- or SDLC-framed. Although this represents a vast majority of the traffic in today's SNA, X.25 and LAN router networks, it does not cover the older byte-oriented protocols such as DEC's DDCMP or IBM's Bisync protocols. It's not that these protocols are not compressible or that the typical data is not compressible. The issue is that these protocols are generally half-duplex polling protocols. Such protocols are more performance-sensitive to end-to-end transmission delay as opposed to absolute data link throughput. For this reason, these protocols would not benefit as much from the enhanced bandwidth provided by SDC due to the inherent delay introduced by SDC. In addition, these protocols are often deployed on multipoint networks for which SDC is not applicable.

It is important to note that data terminal equipment (DTE) supporting these protocols often is connected to the central site via intelligent concentrators that perform protocol "spoofing" and/or protocol conversion. Intelligent nodal concentrators support many of these protocols but transport the user data on the network links between the nodes and the central site using HDLC/ SDLC protocols. These concentrators are well-suited to SDC environments and eliminate the effects of the older protocols.

Other types of synchronous traffic not suitable for SDC are those synchronous data streams that either are not compressible or do not have recognizable HDLC/SDLC frame boundaries. Examples of this class of synchronous data include encryp-ted data, time division multiplexed (TDM) data, and video and voice traffic.

The applications/environments outlined above are best served with standard, clear synchronous analog or digital transmission services.

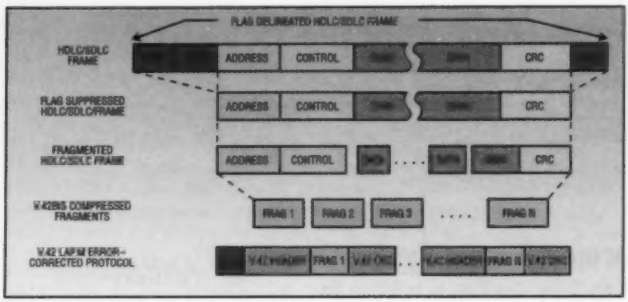

Fig 3 SDC frame compression uses a variety of technologies including error correction and fragmentation algorithms to provide error-free compressed data transmission.

Applications Where SDC Shines

The largest data networking technologies in use today are IBM's SNA, CCITT X.25 and proprietary LAN routers. These networking environments have a common set of characteristics that makes them well-suited to both dedicated or dial-up application of SDC technology. For example:

- All utilize HDLC/SDLC framed data link layer protocols.
- All are predominantly implemented in point-to-point networking topologies.
- All predominantly serve text- and graphics-based business applications for which the data is generally compressible.
- All have a growing need for dedicated and on-demand bandwidth.
- All generally serve mission-critical network applications where back-up and disaster recovery are crucial.

Another, albeit more mature, data networking technology with these basic characteristics is statistical multiplexer networks. Prior to the advent of the X.25 standard and the development of router technology, statistical multiplexers provided the most economical means of concentrating asynchronous and/or mutli-protocol traffic over a high speed transmission media to the central site. Statistical multiplexers are still the concentrator of choice for some asynchronous terminal access networks. The majority of stat mux devices that serve this market utilize proprietary HDLC-like protocols to transport the statistically multiplexed data between nodes. Since the HDLC network channel data in this application environment is statistically multiplexed (not time division multiplexed), it is also an excellent applications environment for SDC-based transmission products.

SDC technology is ideally suited to serve all the networking environments mentioned above. It brings economy and reliability to applications such as the following:

- LAN Interconnect;
- FEP to FEP communications;
- Batch mode file transfer;
- Router link back-up and band width on demand;
- Dial-up routing;
- SNA and X.25 link back-up;
- Dial-up SNA;
- EDI file transfer;
- Disaster recovery;

Performance Expectations And Influencing Factors

So what can one expect from SDC in terms of performance? And exactly what savings can be gained from its use when compared to the alternative digital services that offer similar throughput?

Link performance is best expressed in terms of effective throughput in bits per second. Throughput on a SDC link is dependent on several factors: compressibility of the data, operating speed of the transmission "data pump," link error rate and retransmission rate, configured DTE interface speed, configured HDLC/SDLC link protocol window size, configured HDLC/SDLC link protocol frame size, and higher layer protocol window size.

Compressibility Of The Data

Compressibility is dependent on the amount of redundancy that exists in the data stream. In addi-

tion to text-based information, a wide variety of data types may be statistically multiplexed or packet multiplexed on a single aggregate SDC link. The compressibility of these data types and the aggregate SDC link may vary widely. There are some data types, such as pre-compressed data or binary image data, that can cause expansion when applied to SDC, but algorithms exist that guard against this.

It is important to note that the compressibility of the data is not generally affected by the HDLC/SDLC frame size. That's because SDC continuously updates and maintains the compression vocabulary, so only the similarity in the information redundancy among the frames affects compressibility. In all cases, SDC performs idle suppression on these HDLC/SDLC links. Idle flags are not transmitted, so the presence of any idle time between frames further improves the compression ratios of the aggregate links.

Data Pump Rate

The "data pump" is the basic transmission element of any modem or DSU/CSU. The data-pump rate is the speed at which the base link is transmitting the compressed data. The data-pump transmission speed is applied, therefore, to the compression ratio in order to estimate the actual throughput of the link. If not for the other factors listed above, such as link error rate and retransmission rate, the link's actual throughput would be a simple multiplication of these two factors. In fact, if the other link-performance parameters are configured properly and the link is operating error-free, the throughput should closely approach this figure.

Another important parameter that can affect the throughput of an SDC link is the speed at which the DTE interface is configured. It is important to insure that the attached DTE or nodal/router equipment is configured to transmit and receive data at a sufficiently high rate. Other factors that can introduce throughput limitations on an SDC link are the protocol configuration parameters of the attached DTE or nodal/router equipment. Recognizing that the attached equipment is running a protocol stack that includes, at minimum, a data link layer protocol, there is potential throughput limitation if the protocol parameters are not optimized.

David McNamara *is director of transmission product marketing for Motorola Codex (Mansfield, MA).*

For more info circle 201.

Star Gate and
IBM LAN Distance™

We deliver complete solutions.

Remote network access has never been easier or faster. The best remote access software is now bundled with the best I/O hardware for remote access. All available from a single source–Star Gate Technologies, Inc.

Star Gate is now bundling IBM LAN Distance software with our highest performing I/O controllers–specially optimized for remote access. Supporting from 8 to 64 remote connections at a time, these hardware/software bundles include Standard connection panels for external modem support or Multi-Modem panels with built-in 14.4Kbps modems. No dedicated server required!

Remote Access

To find out more, give us a call at **1-800-782-7428.** And see how Star Gate can simplify your remote access decisions.

Ask for your free copy of Star Gate's "Planning Ahead *for Remote Access,*" a comprehensive overview of remote access technologies.

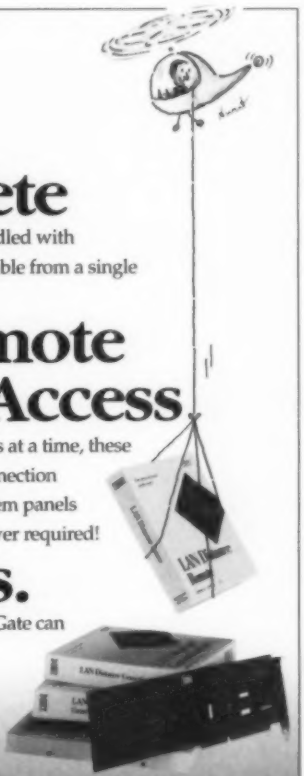

STAR GATE

Star Gate Technologies, Inc. • 29300 Aurora Road
Solon, OH 44139 • 216-349-1860
FAX: 216-349-2056 • BBS: 216-349-2904
INTERNET: custsvc @ sgtech. com

Circle 8 on Reader Inquiry Card

And you can come right along with us. As a local distributor of Upsonic's new A-Series line of uninterruptible power systems (UPS), you can take advantage of a tremendous profit potential. And become an expert on our broad line of A-Series products (System, LAN and Station models), which bring continuous quality power to today's PCs, networks and minis.

Watch your UPS profits soar in '94.

When you join the Upsonic distributor team, you can buy our UPS products at a significantly lower price than you now pay for comparable products. All of which adds up to greater profits for you. If you want to join us as we move up in the UPS market, call Upsonic today.

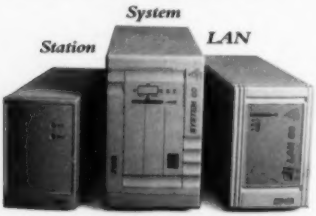

Upsonic's new A-Series products deliver unstoppable power to prevent down time and data loss in today's mission-critical computers.

UPSONIC

29 Journey ▲ Aliso Viejo, CA 92656
TEL (714) 448-9500 ▲ FAX (714) 448-9555
(800) UPSONIC

UPS TODAY

Look where Upsonic's going in 1994!

Ferroresonant UPSs Face PFC Problems

by Aaron L. Davis
American Power Conversion

Recent improvements in the design of computer power supplies meant to end harmonic-induced problems have in turn created a serious incompatibility with ferroresonnant-based power protection. The unexpected problem recently came to light when two major international computer suppliers experienced problems using Ferroresonant UPSs with new Power-Factor Corrected (PFC) power supplies.

How A Ferro UPS Works

In the standby-ferro type of UPS system, the backup power is supplied from an inverter which remains in a standby (off) state until the input AC power fails. When input power is adequate, the power is sent through an electronic transfer switch and then to the load via a ferroresonant regulating transformer. This transformer also has a second input, which is driven by the backup power inverter when the input AC line fails.

In operation, the inverter is normally off and the input power passes through the transfer switch and transformer to the load. When the input power fails, a detector circuit recognizes this event and starts the backup inverter and opens the transfer switch. The load still receives backup power from the transformer; however, the source feeding the transformer has been transferred from the AC utility power input to the backup inverter. Therefore, the principle of operation of a standby-ferro type of UPS is identical to the principle of a standby type UPS except for the use of the ferroresonant transformer.

The ferro transformer can regulate the UPS output voltage for small, slow variations in the AC line voltage when the UPS is in the standby (normal) mode. These variations are passed to the load by most standby type UPS. This is an advantage for use with certain computers using linear type power supplies as it reduces the thermal stresses on the power supply components. However, since the mid-1960s all computers and peripherals use switching type power supplies, which are designed for and totally unaffected by these small, slow variations.

Thus the original design advantage of the ferroresonant UPS has been made obsolete by, and even incompatible with, the advent of modern power supplies in newer computers.

What's The Problem?

Virtually all modern computers incorporate a Switched-Mode Power Supply (SMPS) for converting AC power to the DC voltages required by the computer circuits. After about 25 years without a change, the SMPS is currently undergoing a major design update. This is the inclusion of power factor correction into the basic SMPS design.

The conventional SMPS exhibits an extremely nonlinear load characteristic (low power factor) which causes the computer to draw harmonic currents from the AC power line. In buildings with large quantities of computers installed, these harmonic currents have been rightly blamed for causing fire hazards by overloading transformers and overheating service wiring.

In the European Economic Community, upcoming regulations will effectively ban the conventional SMPS by placing strict limits on harmonic current drawn by computer equipment. User requests and the EEC regulations are forcing computer manufacturers to redesign the SMPS to reduce harmonic currents.

For computers with power requirements above about 300 VA, vendors are incorporating high frequency switching circuits into the SMPS which shape the input current to reduce harmonics. This approach is called "active" PFC.

Most computer vendors already offer one or more products that incorporate active PFC, and virtually all new designs under development above 300 VA have active PFC. This includes file and application servers, disk arrays, and minicomputers.

Why The Incompatibility?

In active PFC supplies the magnitude of the AC input current is controlled by a slow voltage control feedback loop inside the power supply. The characteristic crossover point of this feedback loop is in the region from 4Hz to 30Hz and the crossover is accomplished with less than 90 degrees of phase margin. The result is that the overall system exhibits complex poles with a damped resonance in the 4Hz to 30Hz region. The fact that the power supply is attempting to supply a fixed DC load causes the system to exhibit a transfer function zero in the right half of the complex plane.

Under normal conditions with a low impedance source such as the utility line, the damped resonance of the PFC supply remains controlled. However, when the source does not have low impedance, loading on

the PFC supply will cause the location of the complex poles to move.

With a resistive AC source, the complex poles in the PFC supply move in a direction that changes their frequency, but does not significantly affect their damping or stability. With a more complex source such as a ferro-transformer, the poles move in a direction that ultimately leads to instability.

"The fact that different PFC supplies become unstable under different conditions is due to the variations of the damping factor of the complex poles in the PFC supply. This damping factor is determined by the phase margin of the voltage feedback loop."

The fact that different PFC supplies become unstable under different conditions is due to the variations of the damping factor of the complex poles in the PFC supply. This damping factor is determined by the phase margin of the voltage feedback loop.

A US federal standard, FIPS PUB 94, entitled "Guideline on Electrical Power for Automatic Data Processing Installations" contains a section on the ferro approach to power protection. Some makers of ferro based UPS products selectively quote this standard, failing to mention a number of important sections in which the standard addresses the problems with the ferro design, most notably the potential system instabilities which may result from the unusual output impedance of the ferro:

"If a ferroresonant line voltage regulator were to be placed ahead of another similar unit within an Automated Data Processing product, the combination of the external and internal regulators could become unstable and oscillate."

This standard is available from the National Technical Information Service of the United States Department of Commerce.

Tests were conducted at the request of two major computer manufacturers to determine the extent of the incompatibility. These tests included two standby-ferro type UPSs rated at 500 VA and 5300 VA. The PFC computer power supplies loads were made by IBM, Compaq, and Qualidyne. Variable DC loads were attached to the 5V output of the power supplies. The DC loads were varied from zero to the rated power.

During the tests, the UPS was supplied with 120VAC power, that is, the input AC line was present. The inverter circuits in the UPS were not active and the UPS acted as a simple Ferroresonant transformer. This is the normal operating condition that a user would experience when there is no power failure.

The first tests were done with PFC power supply #1. On the 500VA UPS, the output became unstable when the load drawn by the power supply was adjusted above 90 VA. On the 5300 VA UPS, the output remained stable; however, the power supply was only operated up to its limit of 350 VA which is a small fraction of the capacity of the UPS. With power supply #2 on the 500 VA UPS, instability was found when the load exceeded 25 VA. On the 5300 VA UPS, the output became unstable when the power supply was operated above 450 VA. When the power supply was loaded to 600 VA on the 5300 VA UPS, the voltage and current waveforms were recorded. In this case the voltage at the output of the UPS went far above the normal operating range of the UPS and exceeded 150V RMS during the oscillatory peaks.

Power supply #3 was operated to its capacity limit of about 400 VA without incident on both the 500 VA and the 5300 VA standby-ferro UPS.

In order to determine if the instability was related only to this brand of UPS, a 1000 VA model of a second brand of Ferro UPS was tested and the instability was also found.

The experiments demonstrated that recently designed Power Factor Corrected (PFC) computer power supplies can become unstable when used in conjunction with ferroresonant equipment. This shows up as an uncontrolled oscillation of the AC voltage supply to the computer. The PFC supply and the standby-ferro UPS become unstable under some typical operating conditions.

In addition:

■The stability of the system is reduced when the DC load on the PFC supply is increased. However, the instability did not require that the UPS be in the inverter mode, only that the UPS be operating as a Ferroresonant regulating transformer.

■Different PFC supplies reach instability at different loading points when fed from the same Ferroresonant source.

■Oversizing the UPS for the load will help to reduce the chance of instability. However, in some cases the UPS may need be oversized by a factor of 10 with respect to the load. Therefore, oversizing to achieve stability is not a cost effective option.

■Ferro-equipment can become unstable when loaded

with certain modern PFC computer power supplies. This instability can cause the computer to shut down or be subjected to wild AC voltage swings outside of the safe operating range of the computer. Obviously, any power protection equipment which causes unsafe operation must be considered incompatible with a PFC-equipped computer. These lab experiments do not show that all Ferro-based power conditioning will be unstable with PFC supplies, or that this is the only type of equipment which may exhibit incompatibility. However, user complaints and lab experiments have shown that ferro-equipment is much more likely to cause an instability.

Can Ferros Be Fixed?

While currently small, the number of PFC-equipped computers is expected to grow very rapidly as new legislation take effect and computer manufacturers release new models incorporating PFC power supplies. Since there is a large installed base of PFC incompatible power protection equipment, the number of incidents of PFC instability and resulting down-time will also grow dramatically.

Users may rightly ask whether the problem can be eliminated, and further experimentation has determined that adding a resistive load to the standby-ferro UPS may partially answer the problem. Unfortunately, this approach has its own drawbacks. The customer now gets reduced UPS VA capacity while still paying the same price, and since non-standard modifications have been made, reliability may be affected. Moreover, performing such a fix for an extended installed base of customers would appear to be a daunting task.

According to *Reliability Ratings*, one standby-ferro UPS manufacturer has claimed to have made design modifications to its products, without disclosing what those changes have been.

For the time being, the best course of action for users and system integrators may be, as *Reliability Ratings* recommends for those considering a standby-ferro UPS, to "get written compatibility guarantees from [the] vendor..."

Aaron Davis *is director of corporate communications for American Power Conversion (West Kingston, RI). He holds a B.A. from the Johns Hopkins University and a M.S. from the Columbia University Graduate School of Journalism.*

For more information circle 205.

Just think of it this way... No Power, No Network, No Work, Period.

Why risk it?

Depend on Deltec's PowerRite Pro Uninterruptible Power Systems for unmatched LAN protection. From file servers to workstations to nodes — the PowerRite Pro with Deltec's patent pending Advanced Battery Management (ABM™) will protect your valuable network for up to **10 YEARS!** And, starting at just $449, you can't risk the chance of a downed network due to faulty power. But the features don't stop there — just take a look and compare:

Feature Comparison

	Advanced Battery Management	Triple Protection Warranty	Serial RS-232 Diagnostic Battery Check	High Voltage Regulation
DELTEC	YES	10 Years	Yes	Yes
APC	No	1 Year	No	No

Call our networking experts today at 1-800-DELTEC-1 for more information on our award-winning PowerRite Pro.

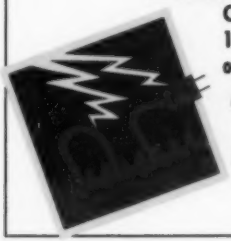

Ask about our VAR Program!

yes NetWare Tested and Approved

PC MAGAZINE EDITORS' CHOICE September 14, 1993 PowerRite Pro PR660

Circle 38 on Reader Inquiry Card

A Layered Defense Plugs The Virus Security Gap

by Susan Richardson and Ronald Sullivan
Reflex Inc.

In the weeks prior to March 6, 1992, computer users worldwide had one thing on their mind. How severely would the Michelangelo virus strike? Estimates that millions of computers were infected screamed out from all forms of media.

The climax came not with a bang, but with a whimper. A mere 15,000 computers were affected worldwide. Nevertheless, Michelangelo brought the virus problem to the forefront, and in its wake came a barrage of virus technology focused on virus detection—scanning—as the primary anti-virus method. Currently there are several scanners on the market which claim "complete" virus protection. However, how safe is an organization which relies on scanning alone to combat the virus problem?

Two captains at the Air Force Institute of Technology addressed this question in a recent report, How Effective are Anti-Viral Toolkits in Preventing Computer Virus Attacks? (December 1993). The report concluded that detection-focused scanning software falls well short of providing effective virus protection.

Their conclusion was that only a preventative "quarantine approach," which employs a combination of anti-virus software and enforcement software to prevent infection, will provide complete anti-virus protection. In other words, the only effective anti-virus approach is to keep the virus off the PC to begin with.

"Most industry experts agree that there are currently over 2,000 PC viruses, and research suggests that there are between 75 and 100 new viruses introduced each month."

Most industry experts agree that there are currently over 2,000 PC viruses, and research suggests that there are between 75 and 100 new viruses introduced each month. As the virus population increases, so does the probability that your system will become infected.

Price Waterhouse research (Annual IT Review, 1992) showed that 65% of computer installations have suffered some kind of direct or consequential financial loss as a result of a virus prompted system failure. In addition, 7% reported that these system failures were major or crippling in scope. The virus problem is common enough that organizations which do not take effec-

tive preventative measures should expect to eventually experience financial loss from infections.

Anti-Virus Weapons

Anti-virus tool kits typically rely on two primary technologies: digital scanning, behavior blocking, or any combination thereof.

Scanning, or static analysis, is a process of examining storage media, or files, to determine whether or not a known virus is present. Although most scanning software companies claim their scanners are 95% effective, the Air Force Institute of Technology research shows that the effectiveness may be as low as 50%, due to the inability of static analysis to account for all viruses.

Scanners are a critical component in the fight against viruses, but scanning as the only line of anti-virus defense has several inherent limitations: they must be continually updated, as new viruses are continually introduced into the computer environment; they don't catch everything—some "stealth" viruses are designed to hide from scanning software; and they neither stop the spread of viruses nor identify the source of the virus.

At best, scanning only solves one of the problems related to computer viruses: detection. Once a virus has been discovered, the organization must find the source of the virus to prevent re-infection. In addition, the virus must be removed before it can do significant damage, which entails certain costs.

Network Loadable Module (NLM) Based Scanners work in similar fashion to their PC-based counterparts, except they reside on the fileserver, thereby centralizing the anti-virus activities. NLMs have the same short-comings as the PC based scanners, plus they may hinder network performance and can easily by circumvented by viruses or users.

Behavior blocking refers to the use of activity monitoring to detect virus-like activity. Behavior blockers may be disabled by some anti-virus software, or the virus itself.

The Layered Defense

Studies have shown that approximately 90% of virus attacks are transferred via floppy disks. Therefore, effective anti-virus technology should include an approach which will quarantine the floppy disks within the organizations environment. The key aspect to avoiding financial loss from virus infection is preventing the PC from getting infected in the first place.

Effective anti-virus technology should include a layered approach as follows: scanning software for detection & identification; virus prevention software so that once a floppy disk is found to be virus free, a "security shield" can be placed on the diskette to insure it remains virus free; behavior blocking software to stop the spread of new viruses that a scanner may not recognize; and user intervention software to prevent a user from circumventing the protection (e.g. by booting from a floppy disk, removing the anti-virus software from the system, etc.). These four types of software combine to provide an anti-virus security solution that stops the spread of computer viruses while simultaneously providing protection against infection.

"The virus problem is common enough that organizations which do not take effective preventative measures should expect to eventually experience financial loss from infections."

The design of one product, Disknet, illustrates this layered approach to virus security. It is a memory resident utility which acts as a shield against virus infection. Once a disk has been certified as virus free, Disknet "signatures" the disk to be used within the environment. If the disk leaves the environment and is written to, it must be re-certified as virus free before entering the environment again. In addition, Disknet includes a utilities package with both behavior blocker and user intervention software.

Susan Richardson *is chief marketing officer at Reflex Inc., (Brier, WA). Ronald Sullivan is the chief financial officer*

For more information circle 116.

Separating Fact From Fiction In SMS (Storage Management Services)

SMS is a set of specifications designed to solve the problems of managing data in heterogeneous environments.

by Michael Kidd
Palindrome Corp.

There has been a swirl of controversy surrounding the Storage Management Services (SMS) specification in recent months. Users are receiving news, both good and bad, concerning this evolving standard for backing up and storing data. Amid the controversy, a number of fallacies have begun to circulate regarding the specification and its implementation in backup and storage management software packages.

Why SMS?

SMS is an industry standard strategy for simplifying and strengthening backup, archiving and storage management support for current and future NetWare versions. It also provides developers with a standard Application Programming Interface (API) set for accessing data on a variety of computing platforms. This standard interface simplifies support for heterogeneous environments by providing a consistent view of heterogeneous file systems and non-file system based resources such as NetWare Directory Services, SQL databases, the NetWare Bindery and other services.

It is important to note that Storage Management Services itself is not a product. Instead, SMS provides an architecture or framework for the development of storage management engines that facilitate data management and protection of heterogeneous environments.

As an open standard, SMS publishes a set of API modules designed to provide safe, reliable backup and storage management application development in a NetWare environment. These APIs are available to any interested software developer. When using SMS APIs to develop backup and storage management software, a developer can:

- Write software to one platform (SMS);
- Let SMS handle all operating and file system protocols; and
- Provide safe, efficient access to virtually any operating system client.

This industry-standard API set for NetWare provides clean and easy backup and restores for non-NetWare file systems, bypassing NOS-specific protocols, file formats and tape formats.

Without SMS, NetWare backup software must be written to each native file format supported under NetWare. To support attached client workstations or servers, the software must be written to adhere to each client's communications protocol and file system. As a consequence, backup software that is not SMS-compliant must manage a bewildering array of file formats, file systems, operating systems and protocols to insure that the backups and restores are handled safely.

Full SMS compliance also calls for full support for the SIDF (System In-dependent Data Format) tape format. SIDF is an open, non-proprietary tape and media format, allowing media interchange.

What are the SMS specifications? It consists of two main specifications. SIDF specifies a method for "packag-

ASIA-PACIFIC

I.T. TIMES

INCORPORATING COMPUTER TECHNOLOGY REVIEW® USA

Today's Technologies For The Asia-Pacific Systems Integration Marketplace

Asia–Pacific I. T. TIMES, which incorporates COMPUTER TECHNOLOGY REVIEW of the USA, is the first and only IT publication from this region that is dedicated to serve the current technology-driven systems integration marketplace, meeting the information needs of decision-makers who are corporate and technology managers of OEMs, VARs, systems houses and volume users who integrate application-specific value-added systems in 14 countries in the Asia-Pacific region: Singapore, Hong Kong, Taiwan, Korea, Japan, Australia, New Zealand, Philippines, Indonesia, Malaysia, Thailand, India, Pakistan and Sri Lanka.

TECHNOLOGY SELLS! And, a well-planned advertising strategy in Asia-Pacific I. T. TIMES gives you that extra mileage of exposure and publicity that is essential in keeping your company and your products competitive in Asia-Pacific's fastest growing segment of the IT industry in the 1990's — SYSTEMS INTEGRATION.

SYSTEMS INTEGRATION YOUR MARKET TODAY

SINGAPORE: **R M TECHNOLOGY MEDIA PTE LTD**
1 North Bridge Road, #16-01
High Street Centre,
SIngapore 0617.
Tel: 334 0393 Fax: 334 3097

HONG KONG: **ASIA–PACIFIC TECHNOLOGY COMMUNICATIONS LTD**
6th floor, Tung Hip Commercial Building 224-248, Des Voeux Road, Central, Hong Kong.
Tel: 853 7616 Fax: 544 4287

THAILAND: **MAN GROUP CO. LTD**
300/31-32 Rachdapisek Road, Bangkok 10310, Thailand.
Tel: 2777-041/835 Fax: 275 7005

Circle 62 on Reader Inquiry Card

ing" data independent of the host OS, as well as a standard open tape format. The Target Service APIs, along with the standard TSAs (Target Service Agents), give storage management software one consistent interface to all network nodes and file systems. Together, they offer a workable solution to the problems associated with proprietary, vendor-dependent LAN storage management systems. At present, SMS is the only way to reliably back up and restore NetWare 4.01 and NDS (NetWare Directory Service).

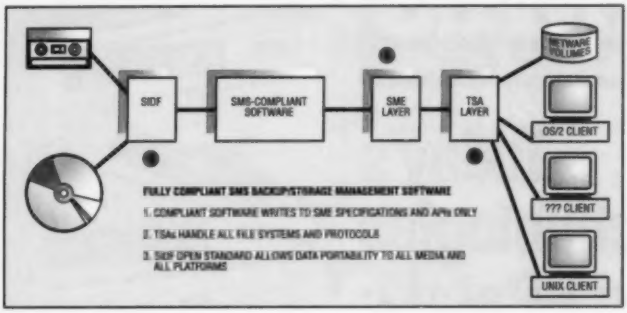
SIDF
SMS-COMPLIANT SOFTWARE
SME LAYER
TSA LAYER
NETWARE VOLUMES
OS/2 CLIENT
??? CLIENT
UNIX CLIENT
FULLY COMPLIANT SMS BACKUP/STORAGE MANAGEMENT SOFTWARE
1. COMPLIANT SOFTWARE WRITES TO SME SPECIFICATIONS AND APIs ONLY
2. TSAs HANDLE ALL FILE SYSTEMS AND PROTOCOLS
3. SIDF OPEN STANDARD ALLOWS DATA PORTABILITY TO ALL MEDIA AND ALL PLATFORMS

Fig Three layers make up the Storage Management Services protocol.

SMS Architecture

The Storage Management Services architecture consists of three primary components (See Fig):

■Storage Management Engine (SME). The heart of the SMS architecture is the Storage Management Engine. The SME communicates with network clients through a common interface that ensures broad platform support and the ability to support new platforms and versions without requiring a change to the SME. The most common implementation of the SME is a backup/restore application. The SME contains most of the value added features and is the component most often provided by third-party developers. Novell has provided a basic SME with NetWare called SBACKUP.

■Target Service Agent (TSA). The Target Service Agent is a software module that isolates the target operating system's data from the functionality of the SME and is the only module in SMS that processes the target's data. A target is any object on the network that requires backup, including: NetWare servers, SQL database engines, print servers, NetWare Directory Service databases, DOS, OS/2, UNIX and Macintosh workstations, etc. All other modules treat the target's data as black boxes . The TSA's job is to prepare the target services' data for backup/ restoration and communicating with the SME. This means that the backup/restore product can support new releases of an operating system as long as there is a TSA that understands the data. With the introduction of NetWare 4.0, Novell has supplied TSAs for NetWare 3.1x servers, NetWare 4.0x servers, NetWare Directory Services (NDS), DOS, Windows and OS/2 clients. Target Service Agents for NetWare 2.x, and Macintosh and UNIX clients are expected to follow shortly.

■System Independent Data Format (SIDF). SIDF is an open standard for the format of data on storage media. In addition, the SMS TSAs generate data streams in the SIDF format within the Storage Management Services architecture. The SIDF specification also defines a standard format for storing data on secondary storage devices such as tape. SIDF is an extensible format that ensures compatibility between future and current SIDF based products and stored (i.e., archived) SIDF data. Novell has formed an SIDF committee to work with industry standards organizations such as American National Standards Institute (ANSI), European Computer Manufacturers Association (ECMA) and International Standards Organization (ISO) to adopt SIDF as an industry standard format.

"SMS is an industry standard strategy for simplifying and strengthening backup, archiving and storage management support for current and future NetWare versions."

In summary, the SMS specifications and APIs offer a development strategy that encourages software developers, committed to fostering industry standards, a means to provide enterprise data management solutions that are safe, reliable and quickly adaptable to changing environments. When new operating systems are implemented, SMS compliant vendors will be the quickest to respond with software support. As networks grow to include clients with a number of network and workstation file systems, only SMS compliant vendors can provide storage solutions which are reliable and proven. In short, full SMS compliance is a must for providing data protection and management to mis-

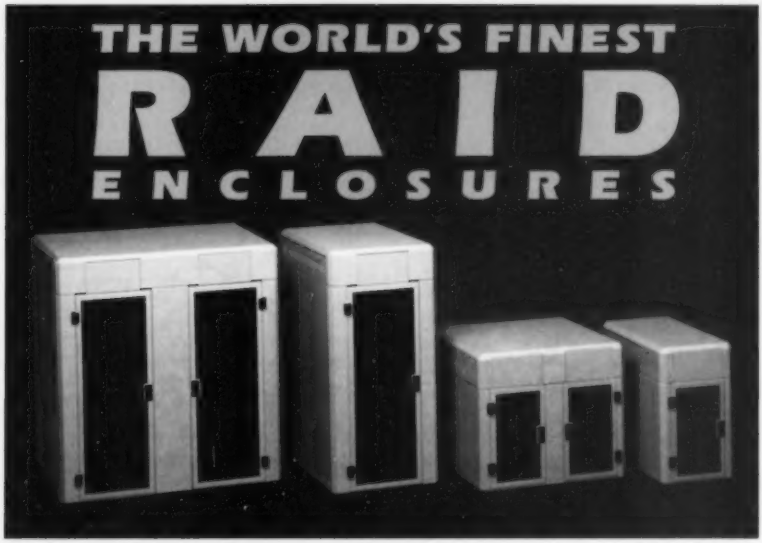

All drive form factors available

Over 125 models of enclosures to choose from

JMR Electronics has been manufacturing quality turnkey peripheral enclosures for more than a decade. We are recognized as a world class leader in the peripheral enclosure industry.

Whether your requirement is for 5¼" full-height, half-height, 3½" (or smaller), removables, CD-ROMs, rackmounts, or RAID towers...

JMR has the solution.

All enclosures are manufactured in the U.S.A. to the highest standards. JMR's complete engineering department, sheet metal, paint shop, mold making, plastic injection molding, cable, and assembly capability make us the fastest and most vertically integrated enclosure manufacturer in the industry.

Whether your need is for one of our standard enclosures or for one that is completely customized, JMR has the capability of producing product expeditiously, and to the highest standards in the entire industry.

Call us today and see for yourself why we have sold more than a million enclosures.

JMR
Electronics
Inc.

Corporate Headquarters.....20400 Plummer Street • Chatsworth, CA 91311 • (818) 993-4801 – Fax (818) 993-9173
Northern California Office.....30 Railroad Avenue • Suite 2 • Point Richmond, CA 94801 • (510) 215-1800 – Fax (510) 215-1820
East Coast Office.....1 Harbour Place • Suite 105 • North Myrtle Beach, SC 29582 • (803) 280-0405 – Fax (803) 280-3082
European Headquarters.....9 Avenue Coghen • 1180 Brussels, Belgium • 322-344-9542 – Fax 322-344-5438
United Kingdom Office.....Unit 9, Chiltern Enterprise Centre • Station Road, Theale • Berks, RG7 4AA • (0734) 304985 – Fax (0734) 304546

Circle 39 on Reader Inquiry Card

sion crit-ical data; anything less is a gamble.

Why is there so much controversy? To support SMS, a vendor providing backup solutions must make a fundamental commitment to supporting open standards. All software calls that access a file system directly need to be replaced by standard SMS calls. The vendor's historically convenient proprietary tape format must be replaced with the standard, extendible SIDF format. As in other parts of the data processing industry, conversion from proprietary solutions to open standards is inconvenient to the vendor, but provides huge benefits to the user.

"To support SMS, a vendor providing backup solutions must make a fundamental commitment to supporting open standards."

When an industry goes through the fundamental change from proprietary to open solutions, there is always a lot of confusion. Any vendor not fully committed to the change will implement "semi-standard" solutions as an interim step, and use any of a number of excuses to explain why a fully compliant solution is not available. SMS is at this threshold today. The standard is complete, the tools are in place, but few vendors have yet delivered, creating a fair amount of confusion in the market.

Fallacy #1: Lots of vendors are fully supporting SMS.

Fact: Novell has defined two levels of SMS compliance. Level One describes products that can use standard TSAs to access network clients. Several vendors have shipped products that can use TSAs, although many of them still make direct calls to the file systems, bypassing the standard when convenient. Proprietary agents do not interoperate with SMS applications.

Level Two compliance requires that the vendor's software not only use TSAs to access clients, but store their data in the SIDF standard format. Most "SMS-based" products on the market today can use TSAs to access data, but lock the data into their proprietary tape formats. To be fair, most have committed to support SIDF at some point in the future, but only as an import/export format—continuing to lock normal backups into the proprietary format.

Fallacy #2: SIDF is a moving target.

Fact: SIDF is a well-defined, extensible tape format. As new file systems evolve, the format will be extended, but not without backward compatibility. Novell is supporting SIDF with SBackup, and Palindrome, Mountain (Filesafe) and Total Network Recall are supporting it as well. Data interchange on tape is now a reality among those vendors that support the standard. As the standard is refined and fully implemented, proprietary tape formats will gradually fall out of favor among software developers—and certainly among end users, who'll demand data portability assured, long-term data accessibility.

"Fact: SMS has been evolving for two years, led by Novell with significant involvement from a number of backup vendors."

Fallacy #3: SMS is a good idea, but it's just too slow.

Fact: When using SMS to backup a server to a DLT tape drive, for example, some developers are seeing sustained throughput in excess of 100 Mbytes/min. This is on real data (not sparse files) without using NetWare 4.0's data compression. As the industry moves to NetWare 4.x, the on-disk compression will make backups even faster, as SMS-based applications can transfer data to tape without decompressing. In these environments, non-SMS applications will actually be much slower.

In order to achieve these high speeds, backup vendors may have to endure the inconvenience of changing some of their file access methods. The performance is there for those willing to make the investment.

Fallacy #4: SMS won't catch on because it requires backup vendors to change.

Fact: Backup as an industry has been among the slowest to adopt standards, but the same forces that have moved other industries to standards-based platforms are at work here. Changing to an open system is never easy or convenient, but it is happening. With swift adoption of the SIDF standard, even reluctant developers may be forced to implement at least partial SIDF capabilities into their products.

Fallacy #5: There's no way to back up NDS (NetWare Directory Services) today.

Fact: A few vendors are already shipping products which fully protect NDS attributes under NetWare 4.x... and the only way they could reli-

ably deliver this backup for NDS was through SMS. By using the SMS-supplied APIs for reaching NDS attributes, those vendors complying with the SMS specification have delivered products which provide full protection of NetWare 4.x. Bindery emulation is not a workable solution for backing up NDS.

Fallacy #6: SMS isn't "ready for prime time."

Fact: SMS has been evolving for two years, led by Novell with significant involvement from a number of backup vendors. The APIs provide a standard, stable and high-performance platform. TSAs available today back up NetWare 3.x (including all name spaces), NetWare 4.01 with NDS, DOS, OS/2 and Windows workstations. In the next few months, TSAs for Macintosh, UNIX NFS and NetWare 2.x will be available. Efforts are underway to provide TSAs for other operating systems and all popular database engines. By the middle of next year, SMS will be the only workable approach to managing storage on complex LANs.

Fallacy #7: SMS is just a NetWare-centric specification.

Fact: SMS and SIDF are specifications originally developed by Novell to solve the growing problems of protecting data in a heterogeneous computing environment. SIDF is now a multi-vendor standard. The SMS APIs are portable to other environments. The SMS specification is portable, the first implementation is on NetWare.

Advantages of fully compliant SMS software are assured compatibility with NetWare; easy, clean migration to heterogeneous environments and additional OS support; data interchange with SIDF, abandoning proprietary media formats; and immediate, certified support for new versions of NetWare and new file systems under NetWare

Disadvantages of SMS 'aware' or 'compatible' software are compatibility with NetWare not guaranteed; heterogeneous and additional OS support depends on proprietary implementation; proprietary media formats lock customers into vendor-specific software solution; and support for new releases is dependent on vendor-supplied updates.

Storage Management Services' goals are to provide software and hardware independence, improve the quality of storage management solutions available to Novell customers and enable the development of storage management applications independent of NetWare version releases. Having the network itself provide transparent access to data ensures a compatible upgrade path as new versions of the network operating system are introduced.

"Does your product use SIDF as its native format? If the answer is "no," then you're locking your backups into a proprietary trap."

The Benefits

SMS is here today. Its benefits—standardized access to clients in heterogeneous environments, faster support for new releases of NetWare and other operating systems, data interchange and long-term access to data based on an open tape format—are available to any vendor willing to make the commitment necessary to provide full support.

If you want to get all the benefits of a fully SMS Level Two compliant system, ask your vendor these questions:

Can your product use TSAs to access servers and clients? Does it use them exclusively, or does it go around the SMS interface with direct file system calls? A completely standards-based solution will use TSAs, and will adhere to the SMS APIs. Direct access to the file or security system raises the possibility of non-standard storage of data, and the necessity for a vendor-supplied update when the target operating system changes.

Does your product read and write SIDF standard tapes? At the very least, import and export capabilities are required for any hope of data interchange.

Does your product use SIDF as its native format? If the answer is "no," then you're locking your backups into a proprietary trap. The only software that will ever be able to restore that data is from your current vendor. Should you change brands, or should they go out of business, your retrieval options are severely limited.

Can I do a complete, SMS based backup and restore of a NetWare 4.01 server, including directory services? If not, the SMS application is not complete.

The LAN backup industry has remained an island of proprietary formats and protocols. Users have had to wait for each vendor to supply proprietary agents to each new client on the LAN. Backup sessions have been locked into non-standard formats, limiting user's choices of future backup and storage management applications. Two years ago, SMS and SIDF were launched to provide a way out of these proprietary, incompatible systems. Today, the standard is complete and standards-based backup systems are here. If reliable backup of heterogeneous LANs, quick support for new OS releases, data interchange and long-term accessibility of data are important to you, SMS Level Two compliant backup systems are the only answer.

Michael Kidd *is the vice president of marketing for Palindrome Corp. (Naperville, IL). He received a B.S. in Computer Science from the University of Wisconsin.*

For more information circle 112.

CD-ROM / IMAGING / MULTIMEDIA

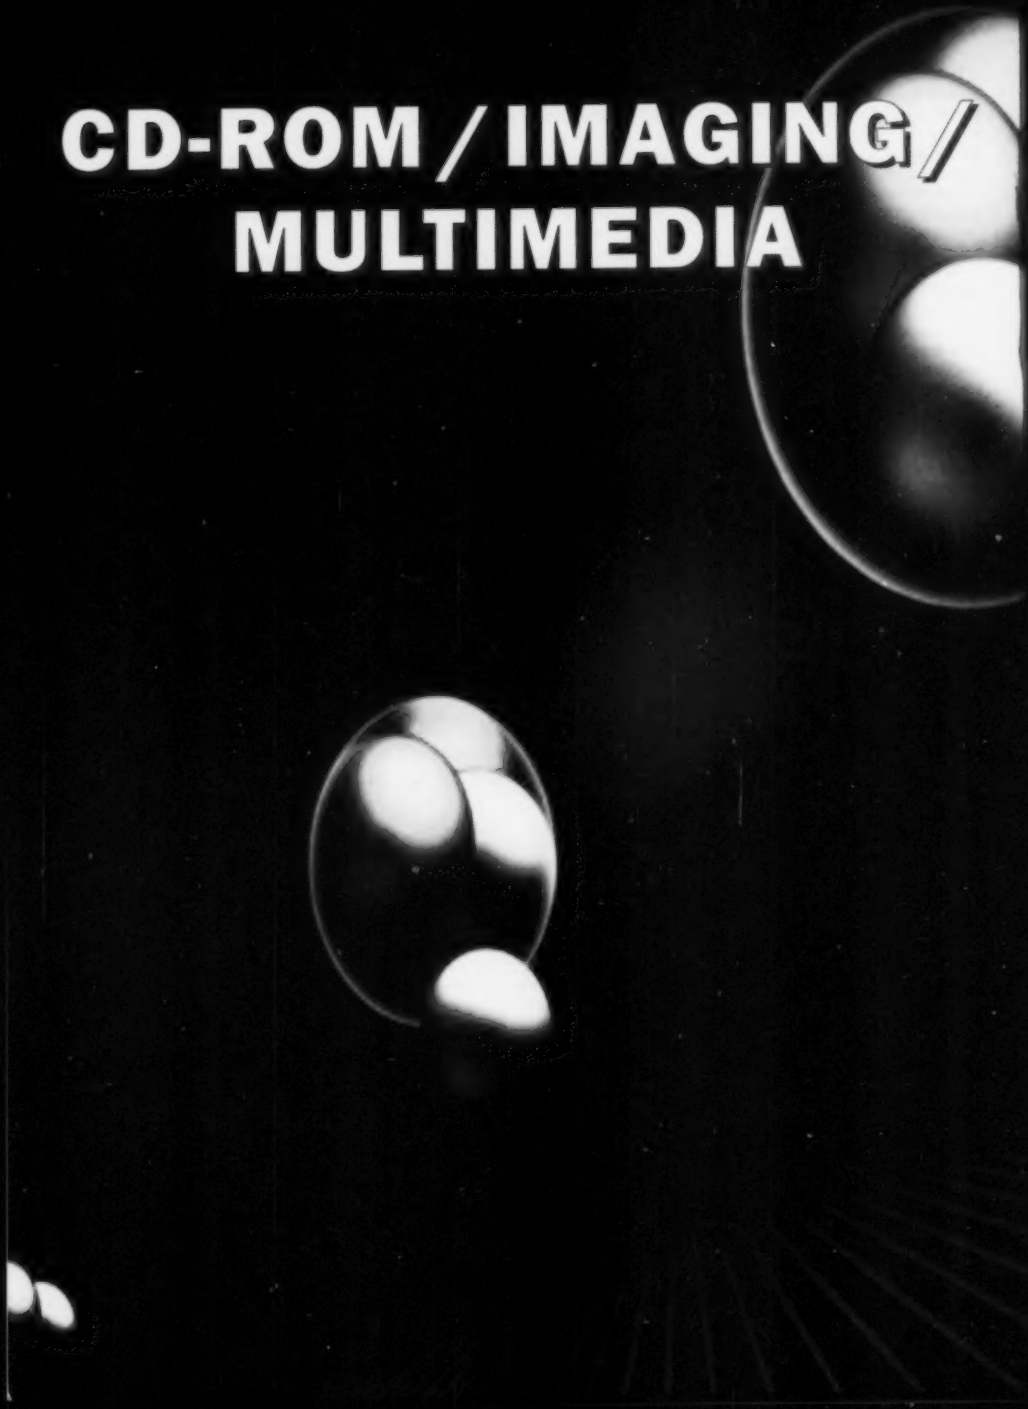

6 CD-ROM Drive

Big $ Smart $

If you need to access data or applications on more than one CD-ROM disc, then Pioneer's automatic changers are unbeatable. Choose from double speed or superfast quad speed changers and save big bucks over multiple single disc drives.

To come close to matching Pioneer changers you would need six individual CD-ROM drives and pay more than $3870*. For less than a third of that cost ($1095)†, get a smart buy on a DRM-602X double speed, and for well less than half the cost ($1495)†, buy the very quick DRM-604X quad speed drive. End the CD-ROM shuffle by automatically switching between discs in seconds and have literally 3,600 MBytes of data and applications at your fingertips.

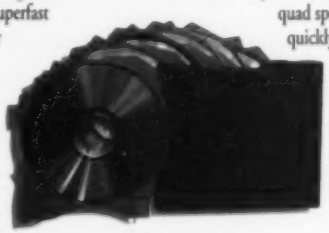

Just load six CD-ROMs into a single pop-in magazine.

For more capacity, daisy chain up to seven drives for 42 discs at your command or buy our clever new DRM-1804X quad speed drive with 18 discs that daisy chain to quickly access a whopping 126 CD-ROM discs.

All Pioneer CD-ROM products feature built-in Kodak Multisession Photo CD™ compatibility and cross platform plug and play for PC's, MAC's and workstations. And our changers feature software keyed security to control disc removal.

To load your favorite discs into a Pioneer changer, just call 1-800-444-OPTI (6784) for the dealer near you.

PIONEER®
The Art of Entertainment

PIONEER NEW MEDIA TECHNOLOGIES, INC.
2265 E. 220th St., Long Beach, CA 90810

Award Winning DRM-604X

Circle 44 on Reader Inquiry Card

*Suggested Retail Price of 6 Chinon double speed CDS-535 internal drives. †Pioneer Suggested Retail Price. Photo CD is a trademark of Eastman Kodak.

High RPMs Add Speed To CD-ROMs And Minichangers

by Richard Millett
Pioneer Technologies

Compact discs offer enormous storage capacity and the lowest cost-per-byte of any rotating storage media. Originally designed for audio data, CDs are not renowned for the speed at which data can be retrieved. While there are a number of possible solutions to this performance dilemma, the most straight-forward approach, engineered by Pioneer, is to spin the disc at four times the normal speed, up to 2100 rpm.

On the surface, this seems a simple solution, but there were several engineering obstacles to overcome in making it a reality. The optical head needed to be completely redesigned with a higher frequency re-sponse to handle the increase in jitter and axial and radial runout. In addition, custom VLSI error detection and correction circuitry were used to ensure uninterrupted data at the 4X sustained transfer rate.

The CD-ROM minichanger system can offer disc speeds of four times the standard, resulting in drives with data transfer rates competitive with all but the fastest magnetic media. For most CD-ROM applications, it is the sustained data transfer rate that is the important number, because most CD-ROM applications are of the "streaming data" type; that is, file sizes tend to be large, and seeks tend to be few.

The obvious solution would be to increase the rotational speed of the platter, thereby pushing data past the head at a significantly increased rate. After all, magnetic disks in hard drives are spun at 3600 and 5400 rpm. The problem is that the data stream resulting from a rotational velocity of up to 2100 rpms pushes existing CD-ROM decoding technology to its outer limits. This is due in part to the overhead associated with a CD-ROM's robust Error Detection and Correction (EDC) scheme.

In magnetic media, if a sector is bad, it is simply marked bad in the FAT and avoided. The surfaces of compact discs are much less susceptible to damage, but they do experience a great deal more handling than hard disks. In fact, users never touch the disks in a Winchester hard disk drive.

"The CD-ROM minichanger system can offer disc speeds of four times the standard, resulting in drives with data transfer rates competitive with all but the fastest magnetic media."

Currently there are no CD-ROM EDC chip sets commercially available that can operate at 4X speeds. Custom VLSI is required to operate at this higher data transfer rate. If a read error does occur, the sector is corrected while maintaining the 614 Kbyte/sec transfer rate.

There are three problems associated with reading data from a CD-ROM, and they roughly correspond to similar problems found in other forms of rotating media: jitter, axial runout and radial runout.

Jitter

Jitter is the tendency of a disc's data channel bit to deviate in position with respect to the optimal read data window. Because of manufacturing quality, jitter is not normally associated with factory prerecorded discs, but is more likely to occur in CD-recordable discs.

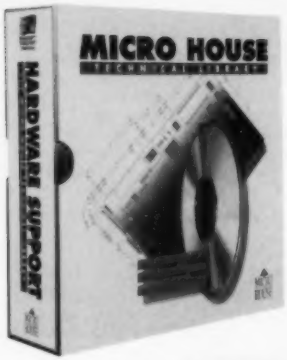

The Micro House Technical Library on CD-ROM contains complete configurations, diagrams, settings, component locations, and vital information on 1500 main boards, more than 700 network interface cards, and 1000 hard drive models. It also contains specifications for an additional 1100 hard drives.

The library contains a full array of specifications from major manufacturers to clone manufacturers including outdated, no longer supported models.

For more details or to immediately start saving valuable support time, call us at **800-926-8299**.

C D - R O M

Hardware Support at Your Fingertips.

MAIN BOARDS

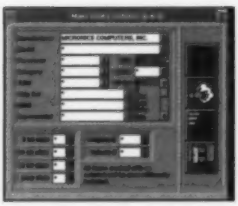

NETWORK INTERFACE CARDS

HARD DRIVES & CONTROLLERS

Quick and easy document retrieval through Microsoft® Windows™.

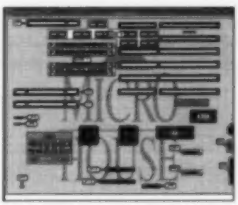

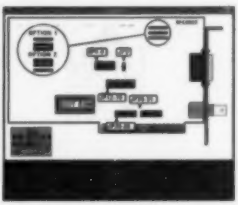

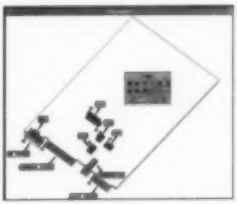

MICRO HOUSE

Micro House International 4900 Pearl East Circle, Suite 101 Boulder, Colorado 80301 vox (800) 926-8299 (303) 443-3388 fax (303) 443-3323

©1993 Micro House International. All Rights Reserved.

Circle 40 on Reader Inquiry Card

In much the same fashion as vinyl records, a factory prerecorded CD master disc is made from the original digital data. Tracks on the CD are spiral, like records, rather than concentric, like those on hard drive disks. Aligned along these tracks are a series of "pits," or channel bits, which modulate the light from the read head and cause the bit patterns. From the master, a mirror-image "daughter" mold is made, and it is from this mold that CDs are pressed.

When the speed of the disc is increased by a factor of four the optimal read data window is decreased by the same factor. Jitter does not affect the head at all, but is detected by the decoding circuitry. If jitter increases beyond the parameters of the decoding circuitry, the track must be re-read, dramatically reducing system performance.

Axial and radial runout are the other two problems that had to be overcome in order to make a drive capable of the 2100 rpm spin rate. Both of these are related to the variation in the position of the data track with respect to the read head of the drive.

Axial Runout

Axial runout, which has a variety of causes, refers to the vertical deviation of the disc as it spins. Simply stated, the rotating disc moves towards and away from the head. Hard drives don't face this problem because the head maintains its position by floating above the disc, on a cushion of air passing between the head and the spinning disc.

The objective lens, or head, in a CD-ROM drive must maintain an exact position with respect to the media in order to correctly focus, track and read the data. To accomplish this, the objective lens must position itself using a knife edge focusing system (See Fig).

The knife edge is generated by dividing the holographic lens element in half. The lens itself is mounted in a system that allows motion both vertically, with respect to the surface of the media, and in a tilting motion with respect to a radius of the platter. This second motion allows the head to traverse several tracks without any motion of the main head assembly, reducing seek time and minimizing axial runout.

Radial Runout

Radial runout refers to the tendency of the data track on the disc to move from side to side with respect to the lens. Radial runout can be caused by a variety of factors, including pressing errors, extreme atmospheric conditions (heat, moisture), and rotational deformation.

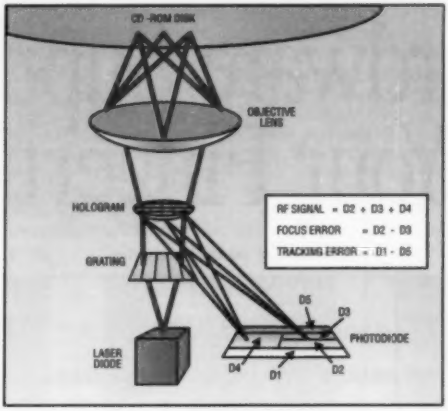

Fig Positioning of the objective lens must be exact in a CD-ROM drive for accurate data tracking and focus.

Accuracy is a primary concern with the increased speed of rotation. The minichanger system uses a three-beam tracking system for accuracy in tracking the data. A low frequency diffraction grating is used to generate the three beams. The brighter zero-order beam is used to read data, while the first-order beams are used as the outrigger beams. These outrigger beams are slightly offset from the center beam and track the edges of the data pits. The drive measures and compares these amplitudes constantly. When the amplitude of the two signals is equal, the lens is on track center.

The ability to track data is due to the Eight to Fourteen Modulation (EFM) code that is used in CD-ROM data. There are no grooves to track on a CD; tracking is done solely on the data pits. EFM is a Run Length Limited (RLL) code based on a (2,10) scheme. A lookup table is used to map 8 data bits into 14 channel bits with three merging bits. This guarantees that there will not be more than three channel bits in a row and there will not be more than 11 blank channel bits. This method provides a constant feedback signal to the tracking circuitry, thus assuring data availability.

Custom servo circuitry for 4X drives has been developed based on video laserdisc technology. In fact, laserdiscs can have a much faster linear velocity across the head than a compact disc spinning at 4X speed.

Pioneer's 4X minichanger can be used in application where CD-ROMs were previously unacceptable, due to either speed or capacity limitations. A six-disc minichanger has a combined capacity of 4 Gbytes, while a jukebox provides a 12-Gbyte capacity. The disc exchange time is less than five seconds. The systems are compatible with all of the existing data standards: CD-ROM XA; Red Book, for audio, Yellow book for CD-ROM data discs; and Orange Book for multi-session CDs.

The adoption of any new technology depends on a number of factors, including cost, stability, maturity and accessibility. However, the primary criterion is that the new technology provide a solution.

Richard Millett *is the engineering manager of the Optical Memory Systems Division of Pioneer New Media Technologies, Inc. (Santa Clara, CA.) He holds a Bachelor of Electrical Engineering and an MSEE, specializing in laser and electro-optics, from the Georgia Institute of Technology.*

For more information circle 206.

Optical Industry Paced By The Meteoric Growth Of CD-ROM Technology

The optical disk drive industry shipped 2.6 million units into data storage markets in 1992 and is projected to deliver more than 4.5 million units in 1993–a 76% rate of growth. Demand is forecast to expand to 11 million units in 1998, according to Robert C. Abraham and Raymond C. Freeman, Jr., authors of the 291-page Freeman report: Optical Data Storage Outlook, 1993. "CD-ROM drives led the industry parade in 1992, accounting for 86% of all optical units shipped for computer data storage and 48% of total revenue."

"Total drive shipments for 1992 were up 119% over 1991 levels, while revenue was up by 49%," reported Freeman. "This difference in rates is due to impressive growth in CD-ROM market in the face of falling prices for all categories of optical products," he added.

According to market analysis, the worldwide market for all types of optical disk drives will approach $2.5 billion at OEM price levels in 1998, a 15% compounded growth rate from the $1 billion value of the 1992 market. Projected revenues for 1998 are divided between rewriteable and multifunction products at $1.2 billion (49% of total revenue); CD-ROM drives at $1 billion (42%), and write-once devices at $222 million (9%). CD-ROM drives will be the predominant class through the forecast period in terms of unit volumes. Rewriteable and multifunction drives will move ahead of CD-ROM in revenue in 1997 and will lead optical revenues thereafter.

PROJECTED WORLDWIDE SHIPMENTS ALL OPTICAL DISK DRIVES (000s)

	1993	1994	1995	1996	1997	1998
Rewriteable	538	736	390	1,112	1,275	1,415
Write-Once	35	36	42	51	64	78
CD-ROM	4,009	5,620	6,880	7,920	8,980	9,910
TOTAL DRIVES	4,582	6,392	7,882	9,083	10,319	11,403
GROWTH FROM PREVIOUS YEAR	76%	40%	23%	15%	14%	11%
INSTALLED BASE AT END OF YEAR	8,834	14,436	20,775	27,269	33,645	39,596

Source: Freeman Association

WITH DCC TECHNOLOGY, EVEN AUDIOPHILES
WILL BE MAKING TRACKS FOR YOUR PRODUCT.
Even the most discerning ears will perk up when they hear your new Digital Compact Cassette systems. DCC delivers unequaled sound quality, track after track.
DCC technology allows you to offer more...
•digital recording •increased shock resistance
•ability to play current audio cassettes •sampling at 32, 44.1, and 48 KHz •digital error correction.
Specify DCC modules from Philips, the inventor of DCC. Philips gives you everything you need for your perfect DCC system: components such as mechanisms, PCB's, IC's and complete functional systems, including hardware and software development tools. Philips DCC modules take your system to audio optimization!
•LDU 1000... Load/record/playback unit
•DIG 1000 Signal Processor... PCB with complete DCC electronic set
•PASC IC's... Complete Precision Adaptive Sub-band Coding IC set
•Thin Film Heads... 9 track record/playback to 18 track playback
Call today and let Philips help you make a quality sound decision!
It's Time
PHILIPS
1-800-235-7373
Outside U.S. FAX 31-40-732277
PHILIPS KEY MODULES GROUP
Circle 45 on Reader Inquiry Card

MPEG Video CD Standard Hastens Digital Video Playback On The PC

The pathway for adoption of digital video on the PC will become very clear in 6-12 months

There is nothing in computing that quite grabs a user's attention like full motion color video playing on a computer screen. Partly it's the color in motion, and partly it's the fascination in seeing familiar television imagery in an unexpected place. Whatever the attraction, re-searchers know that video is far better retained than sounds or words, and software developers are scrambling to incorporate video into games, training materials, presentations and communications tools.

However, the current generation computers were designed to display color bit-map graphics, but not to display full motion 30 frame-per-second video with the accompanying synchronized audio. One way to create a live video display on computer is to retrofit the computer with special hardware that "mixes" an external video source with the signal generated by the computer's built-in display controller, creating "Video in a Window".

Video windowing is a fairly complicated process, wherein the source video signal, in analog form, is decoded, digitized, and stored in a special memory buffer called a frame buffer, and then mixed with the video stream coming from the computer's display controller using a technique called "chroma keying" (Fig 1). In this design, there are two separate video frame buffers—one for the digitized video, and one for the computer's display controller. Depending on where the video is to appear, the corresponding area of the computer's display memory is set to the chroma key color (typically magenta), and when the chroma key samples are detected, video circuitry replaces these samples with real video data from the frame buffer. As the frame buffer is updated at least 30 times per second, and refresh of the video screen occurs at least as often, live video appears on the computer screen.

Because the active video area is typically less than full screen, samples from the captured video are usually dropped in order for the full video picture to fit in a display window. First generation video mixing devices from companies such as Parallax, IBM ("M-Motion") and Chips & Technologies ("PC Video") were only able to drop video samples from the frame buffer in order to fit the full video picture into a window, sometimes leading to jagged lines caused by aliasing. A variety of second generation chips from IBM, Philips, Auravision and others are able to average pixels when scaling in order to remove aliasing effects.

Video windowing is useful, as it brings full motion video to the computer screen, but the source of this video is almost always analog (video tape, laser disk, cameras, TV cable, etc). Although this source material is digitized when temporarily stored in the video mixing frame buffer, a single full screen image requires on the order of 1 megabyte of storage, so one second of video would require transfer and storage of 30 full frames per second. Very few computers have the transfer bandwidth or storage capacity for 30 Mbytes of data per second. If they

by Howard Gordon
Xing Technology

Small Wonders...Now Everyone Can Have One.

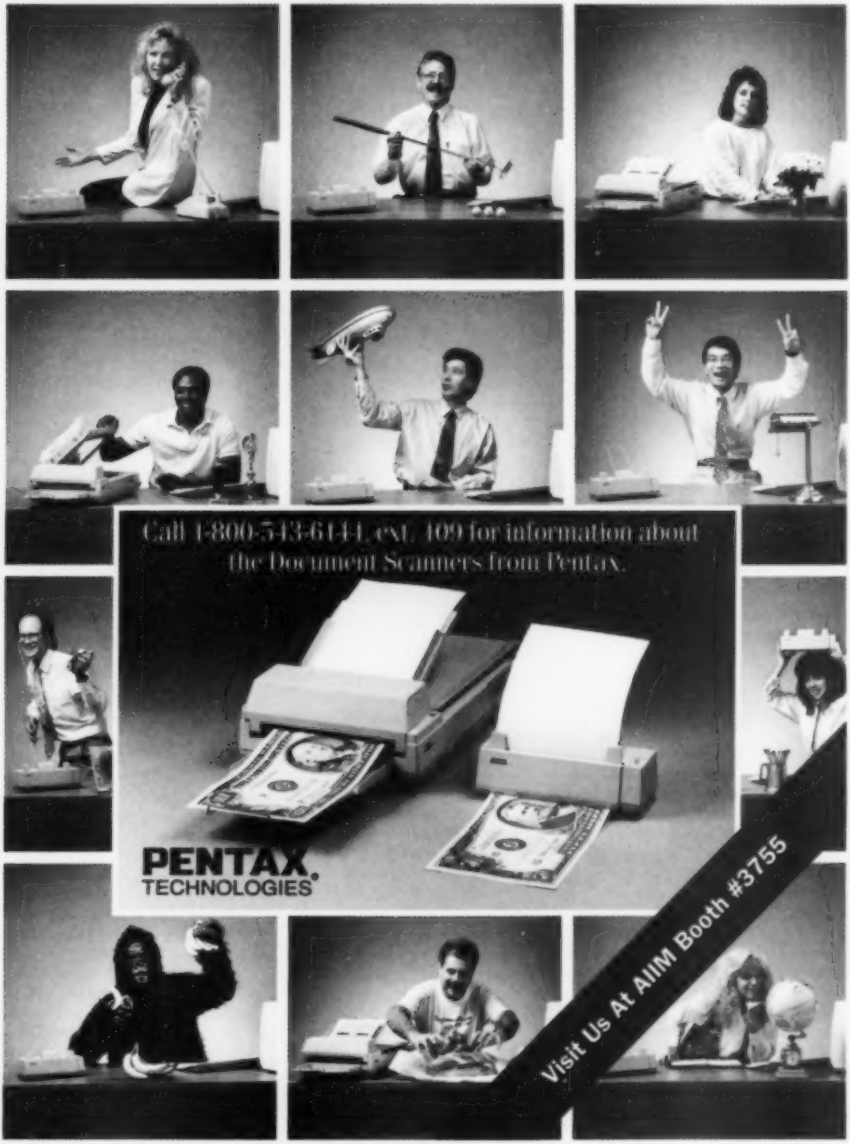
Call 1-800-543-6144, ext. 109 for information about the Document Scanners from Pentax.
PENTAX TECHNOLOGIES
Visit Us At AIIM Booth #3755

100 Technology Drive, Broomfield, CO 80021 Fax (303) 460-1628 BBS (303) 460-1637

Circle 46 on Reader Inquiry Card

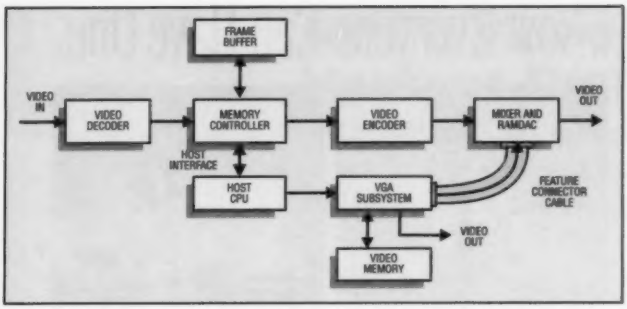

Fig 1 Video windowing is a complicated process, involving multiple video memory buffers, data paths and control devices. Addition of compression/decompression hardware requires further memory buffers, data paths etc.

did, then it would be possible to transfer the digitized video to a storage device, and then play it back later by transferring from the storage device back to the frame buffer or directly to the computer's video memory for direct display.

The Bandwidth Problem

To store the video in digital form, there are several solutions to the bandwidth problem. One possibility is to reduce the resolution of the captured video. The video capture software provided by Apple's Quick-time and Microsoft's Video for Win-dows only transfers a sub-sample of the captured video (typically 1/16 of the actual capture resolution) from a video frame buffer at 15 frames per second.

Unfortunately, that still requires transfer of approximately 1 Mbyte of data per second, significantly exceeding the operating system disk write performance of most disk drives and resulting in very low frame rates for very small pictures. Both Apple and Microsoft provide software to compress the captured video, but only the highest performance systems are capable of absorbing these data rates, and most systems drop 30%-70% of the captured frames.

Another solution is to compress the video data before transferring it to the host computer. There are several existing methods for accomplishing this, including use of JPEG chips from C-Cube, LSI Logic or Zoran to perform 10:1 to 20:1 compression directly from the video frame buffer, Intel's i750 DVI (Indeo) processor for 5:1 to 10:1 compression, or MPEG systems from eidesign, Microboards, C-Cube or Xing Tech-nology to gain 20:1 to 80:1 compression. Some of these systems are quite expensive (the first MPEG encoders cost $250,000), but more recently, live video capture /compression boards have fallen to $500 price levels, with both Apple's Quick-time and Microsoft's Video for Win-dows beginning to support these "real-time" compression systems.

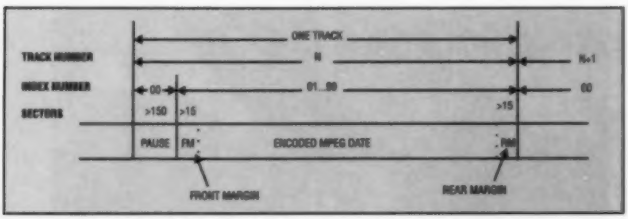

Fig 2 Video CD is a CD-I and CD-ROM XA compatible CD-ROM publishing specification for MPEG audio/video interleaved streams.

Video is generally accompanied by audio. The compression offered by JPEG video or Intel DVI (Indeo) include a variety of techniques for handling audio, but no standard method, while the MPEG international standard fully specifies a nominal 10:1 audio compression format called "MUSICAM", as well as audio/video interleaved format called the "MPEG System Stream". Audio compression usually gets less attention than video compression, because the data rates are significantly lower for audio than for video (stereo 44.1 kHz audio requires 176 Kbytes of storage per second, which is easier to handle than the 30 Mbytes of video), but audio quality has a very large impact on the perceived quality of the accompanying video. Analog Devices (using DSP code from Xing) and Motorola offer "real-time" MPEG audio compression solutions.

Video CD

There is now a CD-I and CD-ROM XA compatible CD-ROM publishing specification for MPEG audio/video interleaved streams called "VideoCD", which was proposed by Philips and JVC, and endorsed by virtually every major consumer electronics equipment manufacturer, including Sony, Matsushita, SGS Thompson, etc. (Fig 2). Paramount Pictures is encoding 50 movies to the MPEG VideoCD CD-ROM format, and other film and television studios as well as game publishers are expected to standardize on VideoCD for content distribution. The VideoCD format will also be basis the for video delivery in video dial-tone and video-on-demand systems.

Once video (and accompanying audio) is stored in a standardized digital format, playback is the mirror image of encoding. The compressed digital video must be decompressed, then deposited into the computer's display controller memory for viewing, or placed in a frame buffer for mixing via a video windowing controller, as described above. Computer systems have the same bandwidth problem with video display that they have with video capture, as display of 30 full screen frames per second requires the generation and transfer of 30 Mbytes of data per second, again exceeding the capabilities of most

current generation computers.

For playback, the compressed video and audio data must first be decompressed, using the computer's CPU or specialized chips that are dedicated to video and audio decompression. Xing Technology is the only company that offers an MPEG software decompression solution for both video and audio, though a number of companies, including C-Cube, SGS Thompson, IIT, Motorola and Philips have developed low-cost video decompression chipsets for compressed MPEG video, and Analog Devices (using Xing DSP code), LSI Logic, Texas Instruments, Motorola and Philips have developed chipsets for MPEG audio decompression (Fig 3).

Somewhere in the system, the video data, now expanded to 20 to 80 times its compressed size, and the audio data, now expanded to 10 times its compressed size, must be temporarily stored before being transferred to the display and audio output devices. Audio storage isn't such a problem, but video storage is, as again there is 30 Mbytes worth of video being generated every second for full screen display. In the case of hardware decompression, this generally requires another frame buffer for temporary storage. In the case of software decompression, there just isn't enough capacity to handle this much data, so video picture size and / or video frame rate must be reduced. Current generation computers can manage transfer rates on the order of 4 to 8 Mbytes of video data per second, allowing one-quarter to one-third screen display with software decompression, but specialized hardware is required to assist the computer in generating the higher data rates for full screen display.

Once the decompressed data has been generated, it must then be transferred to the computer's display subsystem. Current hardware de-compression systems maintain a separate video frame buffer for the decompressed data, and then use the above described chroma-key mixing techniques to display video. This is an inefficient system, as it requires a separate frame buffer for video storage (500 Kbytes to 1 Mbyte of fast memory), complicated and sometimes non-standard cabling (a VESA feature connector cable plus VGA cabling). It would be preferable to transfer the decompressed video directly to the memory of the computer's display controller, but current generation display subsystems are not equipped to absorb the high data rates.

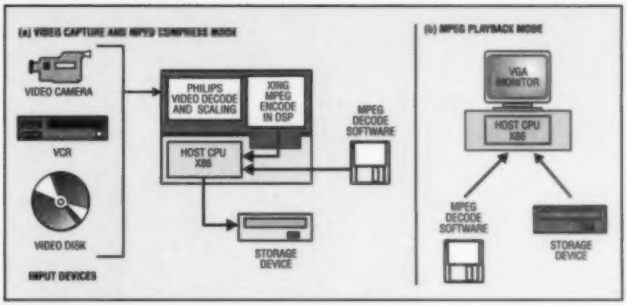

Fig 3 The XingIt! MPEG encoder board operates with any X86 CPU to enable the real-time capture and compression of video signals, as well as MPEG software playback in a window with up to 320 x 240 resolution.

"The VideoCD format will also be basis the for video delivery in video dial-tone and video-on-demand systems."

Next Generation Systems

Next generation display subsystems which can handle the high data rates are now beginning to appear from VGA chipset companies such as Weitek, Tseng Labs, S3 and ATI. New chipsets will accommodate video data with special input ports, storing the video data in the main display memory, and providing special functions such as window positioning and scaling. Ideally, the display chipset would also perform video decompression, and work is under way to accomplish this, now that a publishing standard for compressed video streams—MPEG VideoCD—is gaining popularity. However, no announcements of future devices with built-in MPEG decompression have yet been made, though they are anticipated.

In the future, playback of television quality digital video and audio will be a standard feature of all computing systems, just as the display of color bit-map graphics is now a standard feature of current computing systems. New computer display controller subsystems will support mixing of video and graphics, and in some cases, directly support video de-compression. The MPEG VideoCD standard was an important step in establishing a cross-platform interchange format for compressed digital video and audio, and virtually all providers of video content have adopted MPEG as their format for content distribution. The pathway for adoption of digital video on the PC, once a very confused picture, will become very clear in the next 6-12 months.

Howard Gordon *is founder and president of Xing Technology (Arroyo Grande, CA). He received a Bachelors degree in Mathematics from UCLA.*

For more information circle 117.

Verbatim® tapes, optical and floppy disks. Your best defense against data loss.

Circle 41 on Reader Inquiry Card

OPTICAL STORAGE

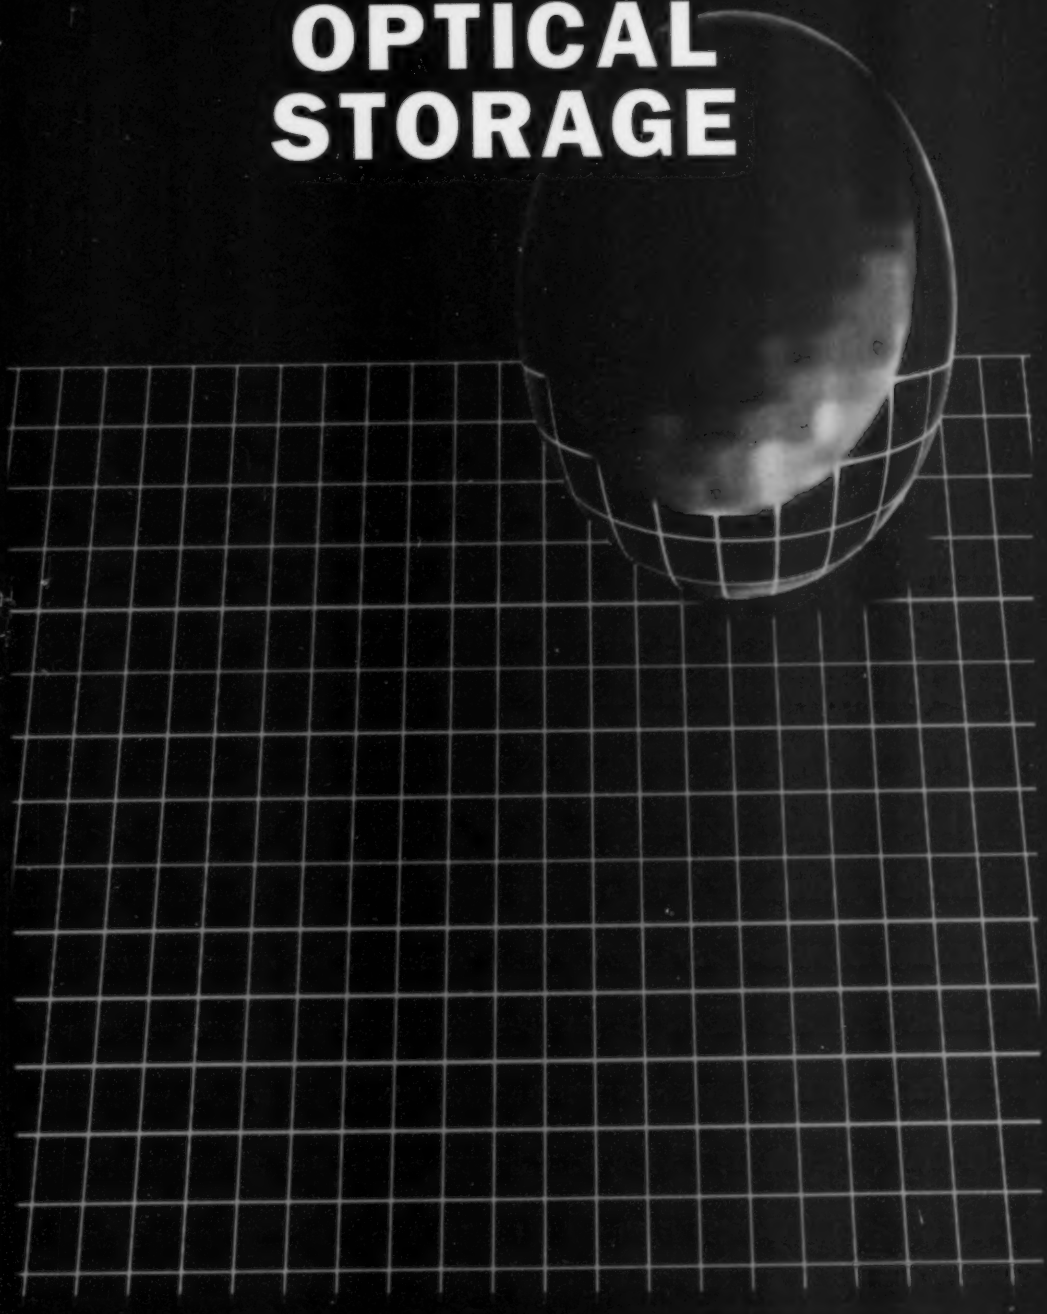

The Easy Road To Information Interchange Optical Standards

The growing list of optical standards takes some of the mystery out of the technology.

by Jeffrey Garner and Arnold Jones
Micro Design International

As large companies get smaller, and small companies struggle to survive, many organizations are looking into the future and planning tomorrow's storage strategy today. While there are many choices for systems that solve current requirements, the big question for end users still remains—How long will our investment in current storage systems last? When faced with the reality that information nearly doubles annually, many organizations realize that today's decision will affect information access and data portability for years to come. When selecting a mass storage solution, users must deal with the dilemma of choosing a system which provides growth capabilities and compatibility features that eliminate the threats common to proprietary systems. In a world where technology marches on, most people are concerned with one major issue: How do we protect our investment? The solution is standardization.

Optical storage devices have gained great popularity because of several key features. The most attractive features include removable media, high capacity storage capabilities, and long term useful media life. Unlike other storage alternatives, optical devices provide users the knowledge that once data is written to an optical disk the information will be safe and secure for more than 30 years. Popular 5.25-in. magneto-optical devices provide 1.3 Gbytes of storage capacity per cartridge, and feature advanced laser technology that eliminates the threat of head crashes found in hard disk systems.

With the growing popularity of optical storage systems, the American National Standards Institute (ANSI) recognized the need for optical standards and established the ANSI X3B11.1 subcommittee to develop a common file system format that would provide universal data access and information interchange. ANSI's X3B11.1 subcommittee made great progress and attracted the attention of the European Computer Manufacturers Association (ECMA) which volunteered the help of their TC15 subcommittee which specialized in file system standards. Under joint development, ANSI's X3B11.1 and ECMA's TC 15 submitted a draft of the file system to ECMA's General Assembly in June 1992. At that time, the proposed standard was accepted, and ECMA recognized the file system as ECMA 167, formally known as "Volume and File Structure for Write-Once and Rewritable Media Using Non-Sequential Recording (NSR) for Information Interchange."

As an internationally recognized ECMA standard, the International Standards Organization (ISO) placed the NSR file system on their fast-track review process. On November 10, 1993, 18 of the participating countries that voted on the ISO standard met in Florida to review several issues of the proposed ISO/IEC DIS 13346 file system. During the November meeting all issues affecting the standard were resolved, making the ISO/IEC 13346 a published file system format which uses the same ECMA naming scheme, "Volume and File Structure for Write-Once and Rewritable Media Using Non-Sequential Recording (NSR) for Information Interchange."

Shift Your Optical Jukebox Into Overdrive By 200-1000%.

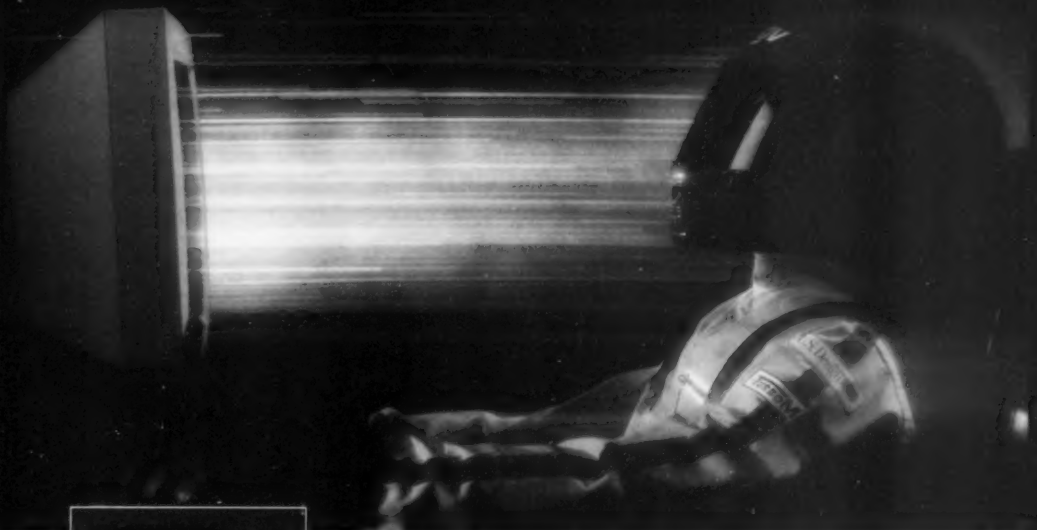

Introducing TurboMASS. It turns your optical jukebox into a lean, mean, storage machine.

Hard-Edged Performance

Want to see an optical jukebox system race? Hook TurboMASS between your system and your jukebox, and watch your back-up, archive and data storage speed up by 200-1,000%.

No More Software

TurboMASS is a revolutionary new hardware and firmware system that eliminates the need for optical software. It serves as a universal translator so that hardware integration from different vendors is simple and cost-effective.

More Reliability

TurboMASS reduces the disk swapping and has a life cycle monitoring feature. It also allows direct interface via RS232 for diagnostics support and even has a built-in battery back-up.

For over 15 years, U.S. Design has been an industry leader in data storage technology. Call us today and find out how the leader just became faster.

U.S. Design

U.S. DESIGN MEANS QUALITY & SERVICE
9075 Guilford Road, Columbia, MD 21046
1-800-622-8732
TEL: 410-381-3000, FAX: 410-381-3235

Circle 80 on Reader Inquiry Card

Now that ISO/IEC 13346 is a globally accepted standard, rewritable and WORM users will benefit by having the ability to share data and media across supported environments and device types (See Table). The goal of this optical standard is to provide the framework for creating file directories, partitions, volumes, data descriptors, along with other attributes that define the file structure which allows universal information interchange.

As a published standard, ISO standards information is available to any vendor that chooses to make their systems compliant with the new file system format. With industry standards each vendor has access to all information dictating the file system format.

"Now that ISO/IEC 13346 is a globally accepted standard, rewritable and WORM users will benefit by having the ability to share data and media across supported environments and device types."

Defining The ISO/IEC 13346 Standard

Although ISO/IEC 13346 is a published standard, developers and end users need to be aware of the elements that make up these file system formats: ISO/IEC 13346 contains five separate parts. The first part contains general references and definitions and the four remaining parts are Volume/Bootblock Recog-nition, Volume/Partition Labeling, File System, and Record Structure. The parts are independent of each other and can be used separately or all together.

TABLE 1: ISO/IEC 13346 FEATURES

- WORM SUPPORT
- REWRITABLE SUPPORT
- VENDOR INDEPENDENT
- INFORMATION INTERCHANGE
- UNLIMITED FILE SYSTEM SIZE
- UNLIMITED VOLUME SIZE
- PUBLISHED STANDARD
- VOLUME SPANNING SUPPORT
- DATA COMPRESSION SUPPORT
- OPERATING SYSTEM INDEPENDENT
- SPARE FILE SUPPORT
- MEDIA DATA PORTABILITY

The volume is a logical address space of sectors on the media, and in regard to optical usually corresponds to one side of a cartridge. Therefore most optical cartridges would be considered two volumes, one for each side of the media. The partition is a contiguous sequence of sectors within a volume which allows a volume to contain multiple partitions. A logical volume is a set of partitions on which a file system may be recorded. A volume set is a set of volumes with a common volume set identification, which may contain multiple logical volumes (file systems). The Figure is an example of a volume set (*research*) consisting of two volumes, *alpha* and *beta*. Volume *alpha* contains two partitions, while volume *beta* contains a single partition. Also shown are two logical volumes, *hardware* and *software*. Logical volume *hardware* consists only of partition one on volume *alpha*. Logical volume *software* consists of partition two on volume *alpha* and partition one on volume *beta*. Logical volume *software* illustrates how a single file system can span multiple physical cartridges.

Part two defines a general method for recognizing what standards were used to record the data on the media. For example this allows a driver to determine if the media contains data structures for ISO 9660, ISO/IEC 13346, or both. Even though the standard supports mixed standards on a single volume, in most cases only one will be found.

The standard also defines a method for recording basic booting information for a variety of hardware systems and operating systems. This is accomplished through what is called a Boot Descriptor. The boot descriptor contains fields that define the system architecture type and an operating system identifier. Multiple boot descriptors may be recorded on a volume, allowing boot support of multiple operating systems for a variety of different hardware systems on a single volume. ISO/IEC 13346 provides a standard method for hardware system boot ROMs to support booting from optical.

Volume/Partition Labeling

Part three of the standard defines a general method for defining volume labels, volume sets, partitions and logical volumes. All this functionality is similar to the DOS utility FDISK which is used to partition a single hard disk for different operating systems. In addition, ISO/IEC 13346 includes the ability to define logical volumes which may span multiple physical volumes common in optical jukeboxes. The logical volume can contain anything, such as a DOS file system, or a UNIX file system, or the NSR file system

"With the growing popularity of optical storage systems, the American National Standards Institute (ANSI) recognized the need for optical standards and established the ANSI X3B11.1 subcommittee."

75 YEARS OF OPTICAL EXPERIENCE WENT INTO OUR NEW 3.5 INCH OPTICAL DRIVE.

It's a claim few companies can make. Yet at Olympus Optical, our unequalled 75-year history of product development in the camera, microscope, medical imaging and consumer audio business taught us what was needed to be the very best in the optical storage business: high quality, high-performance products at a competitive price.

Olympus has become the world's third largest manufacturer of rewritable optical storage devices as an OEM. Now, for the very first time, our 3.5" drives are available to you from Olympus Image Systems. Our new 3.5" MOS 300S subsystem and 300E standalone magneto-optical disk drives include such technological innovations as disk rotation of 3600rpm. Fast data transfer rates (768KB/sec) and seek times (less than 38msec.).

ISO standard 128MB capacity.

And a rugged, reliable, dust-free design. All of which make the reliable MOS 300S and 300E ideal for today's data-intensive storage applications like desktop publishing and CAD applications. What's more, our new 3.5" optical drives are backed by unequalled support programs. And, of course, the proud 75-year heritage of Olympus. Best of all, Olympus Image Systems' optical drives are available now through our nationwide distribution network.

For more information on MOS 300S/E optical disk drives or to find the distributor nearest you, Olympus Image Systems, Inc. invites you to call our Optical Memory Division at 800-347-4027.

OLYMPUS IMAGE SYSTEMS, INC.

Optical Memory Division • 15271 Barranca Parkway • Irvine, CA 92718-2201

©Copyright Olympus Image Systems, Inc. 1993.
All trademarks are trademarks of their respective companies.

Circle 61 on Reader Inquiry Card

Standards

as defined in part four of the ISO/IEC 13346 standard. ISO/IEC 13346 provides a standard method for defining logical volumes for any use, including native file systems or the NSR file system.

Part four of the standard defines how to record a hierarchical file system which has been designed for information interchange. When used in conjunction with part three of the standard, the NSR file system is recorded within a logical volume. Basically, this is the section that defines how files and directories are stored on the media so the device can maximize system capabilities by easily locating information based on data descriptors. File system specifications were written with the goal of data interchange between a wide variety of operating systems. Therefore the standard has the ability to record a full range of file attributes including extended attributes.

Part five of the standard defines record structured files. This section of the standard was mainly written to support backward compatibility to certain features of ISO 9660.

Information Interchange Implementation

Now that you understand what the structure of standardized file systems are all about, there are several key factors that determine if media and data portability will be a practical solution for your environment. There are three main areas that affect information interchange, and each of them need to be considered before the initial investment is made into purchasing a standards compliant optical solution. The areas to consider include: physical level interchange, logical level interchange, and application level interchange.

Physical Level Interchange

Physical media interchange focuses on the actual size of the media supported across different optical drives. The data that is being interchanged at this level is dependent on the ability of different drives to accept media used in another manufacturer's drive. When it is stated that two different optical drives can interchange data at the physical or sector level, it means that the drives can universally read and write data from media between different devices. For example, physical interchange is possible between different 5.25-in.

"There are three main areas that affect information interchange, and each of them need to be considered before the initial investment is made into purchasing a standards compliant optical solution."

optical drives, but not between 5.25-in. and 3.5-in. The physical dimensions are different and therefore not interchangeable.

To ensure physical (sector) level interchange optical media is manufactured to specifications defined by several standards organizations. These organizations include the American National Standards Institute (ANSI), the European Computer Manufactures Association (ECMA), and the International Standards Organization (ISO). Table 2 shows the different types of optical media and the associated physical media standards from the different standards organizations:

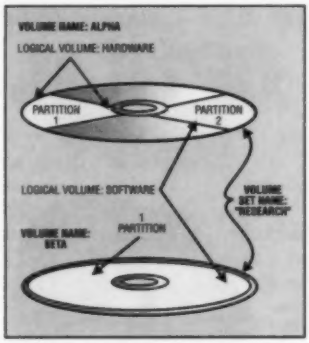

Fig An optical cartridge can be considered two volumes for purposes of capacity planning

Logical Level Interchange

The logical level interchange is sometimes referred to as the file system interchange. This level focuses on the information structure which determines information interchange between different platforms. The logical, or file system structure is concerned with partitions, directory structures, volume information, data descriptors along with other elements that defines how information is arranged on the optical media. Optical media supports a wide variety of native file systems such as FAT (DOS), HPFS (OS/2), HFS (Apple MAC), JFS (AIX), System V (UNIX), among other. With removable optical media and native file system support, data can be interchanged between different platforms running the same operating system. To interchange files between platforms running different operating systems, a standardized file system format must exist to support both operating environments, such as NSR.

Application Level Interchange

Once you have physical (sector) and logical (file) interchange capabilities you can move data between different optical drives on different computers using different operating systems. Application interchange which deals with the contents of the files on the optical media. Lets say that you have sector and file interchange between a PC and a Mac. This tells us that you can copy files to an optical cartridge on the Mac, bring that cartridge over to a PC and type in *dir* and see the files you copied to the cartridge on the Mac. Assume that you used a database application called *alpha* on the Mac to create the files on the cartridge. Now you wish to read those files on the PC using a database application called *beta*. Unless *alpha* and *beta* use the same file format you will not be able to read the data written by *alpha* with the *beta* program. To have application interchange, the applications must be using the same application stan-

dards, and in this case the same file format.

Currently there are numerous file format standards that allow application interchange. There are database file format standards such as DBF and DIF. There are image file format standards such as TIFF, PCX, GIF, and JPEG along with audio standards such as WAV and VOC. Applications that support standard file formats will be able to interchange data at the application level without specialized conversion utilities. What this means to the user is that the file system standard provides access to the data, while the application manages import/export capabilities, thus providing transparent data access, regardless of the operating environment.

ISO/IEC 13346 File System Benefits

Now that we have covered the elements that form the framework supporting information interchange, the real benefits are as follows:

■**Rewritable And WORM Support.** Native file systems, such as FAT, HPFS, JFS, and others, were designed for magnetic media and are able to rewrite a sector on the media. Write-once does not allow sector rewrites, because a sector may only be written to once. Therefore if you wanted to use WORM media it could not be done through the native file system; however, NSR now supports WORM devices. Until NSR was available every WORM implementation used some form of a proprietary file system since this was the only way to support WORM. With proprietary file systems, the ability to have file interchange is lost for those users that need to move data across different file system formats. The NSR standard provides a true standard file system for WORM media that supports information interchange and provides a vehicle for media and data portability.

■**Operating System Independence.** In the same way that CD-ROM media written in the ISO 9660 file system standard can be read on a wide variety of operating systems, media written in the NSR file system standard can be read/written on a wide variety of operating systems that support the NSR standard. As the title of the standard states it was designed for information interchange. For example, you can store files under DOS on a piece of NSR media, then take this media to a system running AIX and update the files, then take the media to a machine running OS/2 and so forth. NSR provides you with file interchange between today's operating systems and the operating systems of tomorrow.

TABLE 2: MEDIA STANDARDS

MEDIA SIZE	STANDARDS ORGANIZATION		
	ANSI	ECMA	ISO
5.25-IN. REWRITABLE-650 MBYTES	X3.212/1992	NONE	ISO/IEC 10089 FORMAT A
5.25-IN. WORM-650 MBYTES	X3.220/1992	ECMA 153	ISO/IEC 11560
5.25-IN. MULTI-FUNCTION-1.3 GBYTES	NONE	ECMA 184	ISO/IEC 13549
3.5-IN. REWRITABLE-128 MBYTES	X3B11-738D	ECMA 154	ISO/IEC 10090

■**Vendor Independence.** Different vendors frequently use different techniques to create application solutions. Implementing unique programming techniques causes problems with file interchange because different vendors may not understand the methodology used to write data, therefore compatibility problems are common. With NSR media you will be able to exchange files with different platforms even if the NSR software on the other computer was developed by a different vendor. Access to your data is not locked to any single vendor, but is vendor independent.

"With NSR media you will be able to exchange files with different platforms even if the NSR software on the other computer was developed by a different vendor."

■**Massive Volume Support.** Most native file systems were designed before the massive capabilities of today's storage subsystems and have limitations such as a maximum of a 2-Gbyte file system size. Unfortunately, native file systems were not designed to support jukebox type storage devices. Unlike native file systems, NSR was designed for the massive storage needs of tomorrow. This flexibility allows you to create volumes as large as you want, with complete flexibility of file system size and volume spanning across multiple platters and devices. By today's standard it is clear that NSR was designed to meet the mass storage needs of a global community well into the future.

To measure how important standards are in the quest for universal data access and investment protection, ask yourself these questions: What operating system will I be using 10 or 20 years from now? What hardware platform will I be on? If I store massive amounts of data today what will I do if I can't read it on my machine of tomorrow?

For most of us the answers will not be too surprising. We need standards compliant systems, and with ECMA and ISO supported devices, users have the assurance of data access for years to come. Both the ECMA and ISO standard file systems provide a high level of confidence in long-term data interchange between today's operating systems and those operating systems of tomorrow.

For further information on optical standards referenced contact the following organizations: ANSI and ISO Documents: Global Engineering Docu-ments 1-303-792-2181; ECMA Docu-ments: European Computer Manufac-turers Association +44-22-735-36-34

Jeffrey Garner *is the product marketing manager and* **Arnold Jones** *is director of research and development with Micro Design International (Winter Park, FL).*

For more information circle 106.

DO YOU HAVE A NEED FOR REMOVABLE STORAGE PRODUCT?

FOR DATA SECURITY, PORTABILITY, BACK-UP OR RAID APPLICATIONS

CALL CRU, INC.
THE U.S. PATENT OWNER OF DATAPORT™

U.S. Patent No. Re. 34,369 and Des. 308,052 for REMOVABLE ADAPTER.

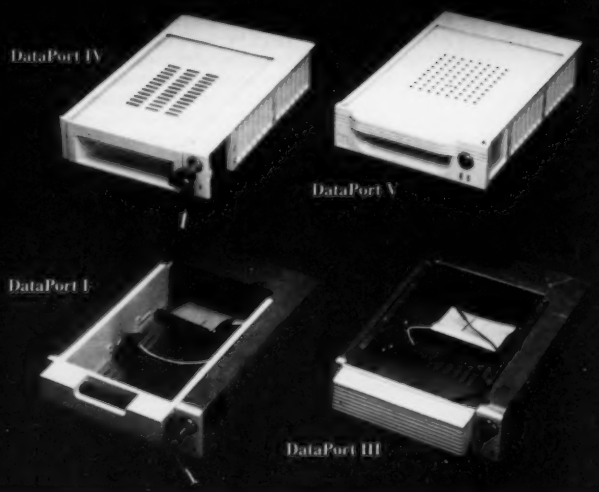

Features
Support most 3.5" Hard drives
IDE, or SCSI Interfaces
Hot Removability
Key lock ON/OFF Security
Fan Cooling Options
Built-in SCSI ID Selector Switch
Aluminum or Low-Cost Plastic Version
Power/Drive Activity LEDs
Patented Protected

CRU offers a full range of RAID and storage enclosures

CRU, Inc. the computer storage experts, offers a full line of patented DataPort removable storage modules (RSM). DataPort is the perfect solution for protecting your valuable data. Its benefits are security, portability, on-line backup, and hot-removability for RAID. DataPort is a lost less expensive than other forms of removable hard drives such as SyQuest, Bernoulli, and Kalok, etc. When used in conjunction with disk mirroring software or a hardware RAID controller, DataPort can be used to quickly change out of a failed hard drive, thereby reducing computer downtime.

The DataPort frame is form-fitted for a standard PC/AT 5.25" half-height bay. Its cartridge can accept most 1" or 1.65" high 3.5" hard drives. Thus, it can easily be installed in minutes. Once installed, the user can remove data simply by turning the key and removing the cartridge from the frame.

Connector Resources Unlimited, Inc. (CRU) has been in the computer storage and peripherials business since 1986. Besides DataPort, CRU offers a full line of storage enclosures such as one bay, two bay, four bay, 3.5" portable and RAID chassis. Call us today at (408) 957-5757 and see how our experienced sales consultants can help you with all your storage needs.

CRU
CONNECTOR RESOURCES UNLIMITED INC.

"simply the best"

Connector Resources Unlimited, Inc.

All Trademarks and Registered Trademakrs are of their respective holders, CRU and DataPort are registered Trademarks of Connector Resources Unlimited, Inc. Disclaimer: All specifications and prices are subject to changes without notice.

Circle 66 on Reader Inquiry Card

1005 Ames Avenue, Milpitas, CA 95035 • TEL: 408-957-5757 • FAX: 408-942-0862

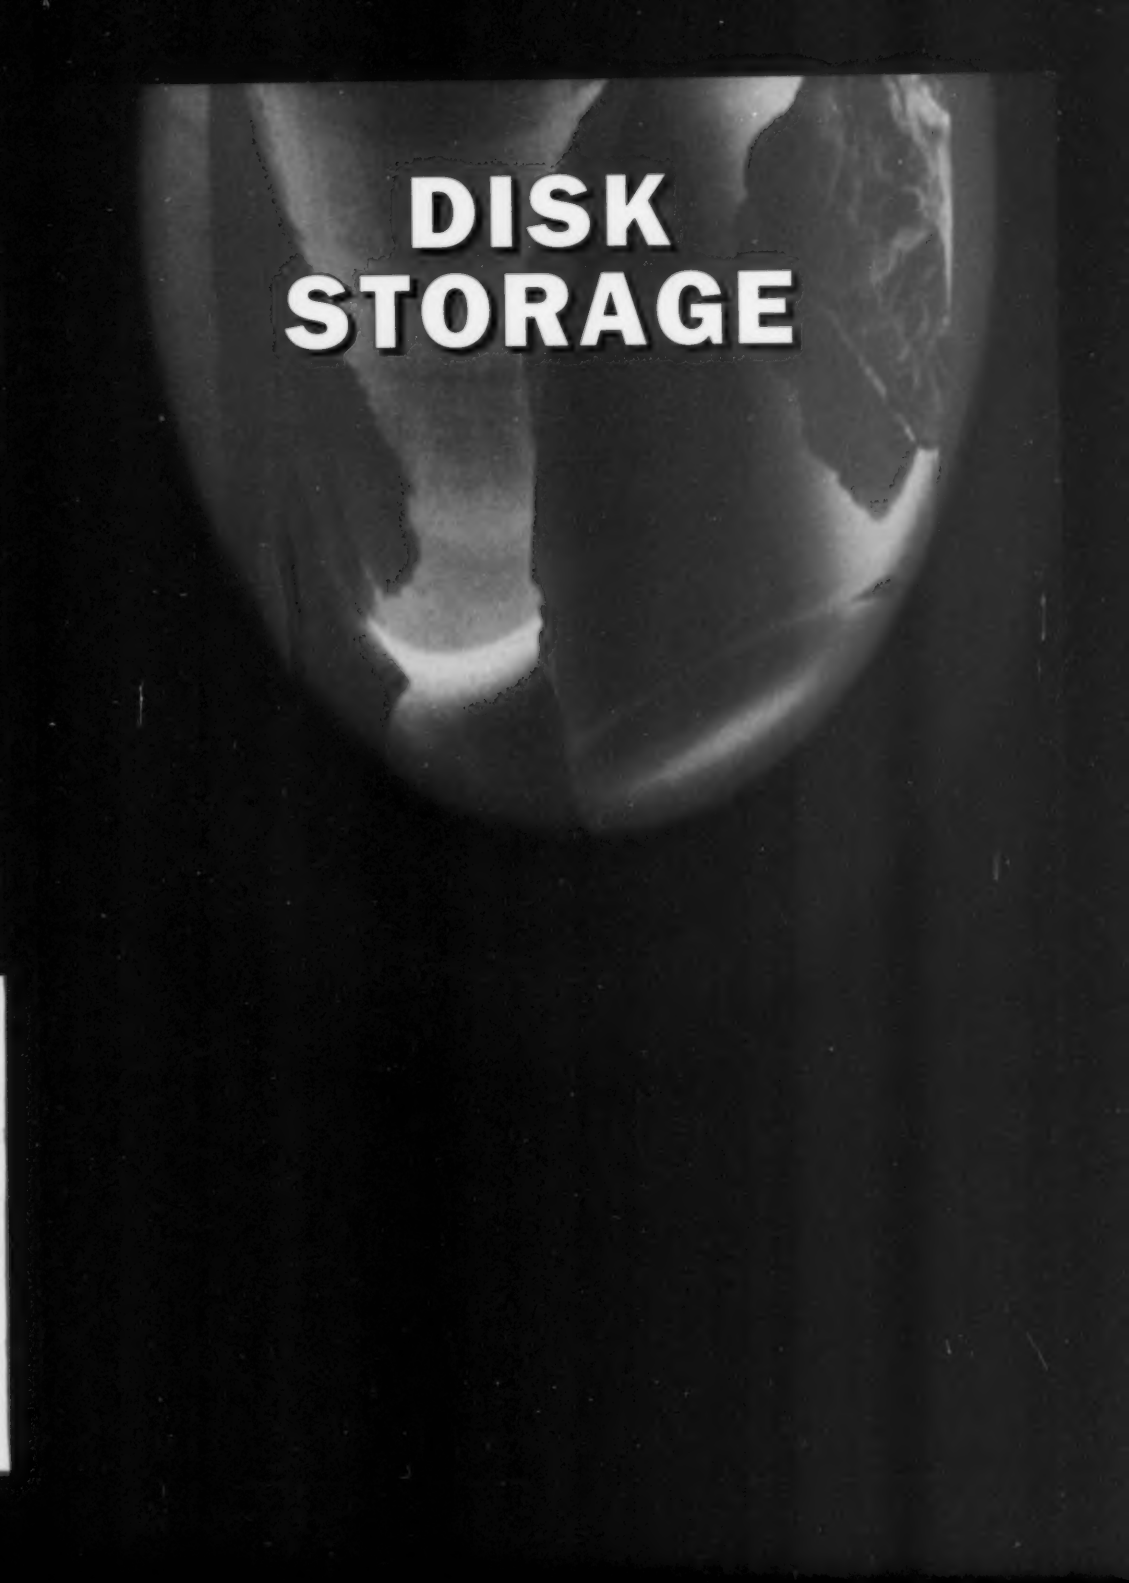
DISK
STORAGE

Flash Memory

Choosing And Integrating Current Flash Solid State Storage Systems

by Markus Levy and Pedro Vargas
Intel Corp.

As the mobile computing trend accelerates in 1994, you will need to configure systems with mass storage solutions that meet new requirements for reduced size and weight as well as increased ruggedness, reliability and battery efficiency. Moreover, your customers may be looking for new ways to bring the productivity benefits of the PC to applications with special requirements for space savings, shock and vibration resistance, and tolerance for extreme temperatures and humidity.

To serve these special mass storage needs, Intel and other vendors are developing a variety of new flash memory-based alternatives to rotating magnetic mass storage. Flash chips, PCMCIA cards, flash drives and peripheral add-in boards are available in a variety of megabyte densities and software configurations.

What Is Flash?

Flash is a semiconductor technology spin-off of familiar EPROM technology. Like EPROMs, flash memory is non-volatile and retains data with-

Introducing The Ultimate

RAM Cards offer from .5 to 2MB of unbeatable storage, with built-in battery backup to protect against data loss.

Flash Cards provide enough capacity (2.5 to 20MB) to run Windows, and feature a low-power design to maximize system battery life.

Meet the most valuable business cards you can carry. Epson PCMCIA cards. Now your notebook computer, palmtop or PDA can have unsurpassed power, storage and communications capabilities, anywhere you take it. Epson cards cost much less, too.

PC Card

All company and/or product names are trademarks and/or registered trademarks of their respective manufacturers. Epson disclaims any and all rights in those marks. Epson is a registered trademark of Seiko Epson Corp.

out requiring a power source. Unlike EPROMs, flash memory is erased and reprogrammed electrically and very quickly. Because of these characteristics, flash can be configured with software to provide a solid-state alternative to disk drives.

Flash And The Traditional PC Memory Hierarchy

The traditional PC memory hierarchy includes ROMs or EPROMs for boot code storage; disk drives for operating system, application program and data file storage; DRAMs for program execution and data manipulation; and floppy disks for file transfer. Today, flash memory is impacting each of these technologies.

Flash replaces ROMs and EPROMs in BIOS applications because it allows the BIOS to be easily updated to fix bugs or add support for new system peripherals. BIOS updates can be shipped to the customer on a floppy disk or over a computer bulletin board. Flash BIOS can give more configuration flexibility and help reduce support costs.

Flash Memory

Random-access NOR flash chips can also be used to store operating systems and application programs, and the CPU can execute these programs directly from flash. In this way, flash reduces the redundancy of the traditional memory hierarchy, where these programs are downloaded from a disk drive or ROM chips into DRAM for execution.

"Flash replaces ROMs and EPROMs in BIOS applications because it allows the BIOS to be easily updated to fix bugs or add support for new system peripherals."

Flash can also be used to configure mass storage devices that take the form of PCMCIA cards or traditional add-in cards. Today's PCMCIA flash storage devices provide from 2 to 40 Mbytes of storage, while add-in boards provide virtually unlimited capacities. Flash storage devices can replace disk drives completely, and removable PCMCIA devices can absorb the functionality of floppy disk drives as well.

As a semiconductor technology with no moving parts, flash provides several advantages over rotating magnetic mass storage. Today, flash costs more on a dollar-per-Mbyte basis, but complete flash storage systems in the 2-10 Mbyte range are available for less than the floor cost of many disk drives. The unique advantages of flash make the technology more attractive in many applications despite the higher price tag.

■ **Size and weight.** Today, up to 2 Mbytes of flash storage can be

Business Cards from Epson.

FAX/Modem Cards come in 2.4 to 14.4 Kbps speeds, work with the widest range of portables, and include award-winning software to keep you well-connected.

Our 20MB Flash Card delivers the ultimate in storage power, with pre-installed compression software to double capacity.

And they're supported by Epson's two-year warranty and free help line, the Epson Connection. For details on specs and compatibility, call **1-800-BUY-EPSON** (800-289-3776). And put more power into your portable.

EPSON®

© 1993 Epson America, Inc., 20770 Madrona Ave., Torrance, CA 90503. For dealer referral or warranty details, call 800-BUY-EPSON (800-289-3776). In Canada, call (800) GO-EPSON. For Latin America, 305-265-0892.

Circle 16 on Reader Inquiry Card

Flash Memory

placed in a single chip measuring 1.2 mm high and 14x20 mm wide. Up to 40 Mbytes of flash can be squeezed into a PCMCIA Type 1 card that weighs less than 2 ounces.

■**Ruggedness and reliability.** PCMCIA flash cards and drives are designed to withstand 1,000 Gs of operational shock, which is five times higher than the most rugged disk drives. Flash cards and drives also operate in environments producing 15 Gs of omnidirectional vibration. Temperatures ranging from 0° to

"Flash storage devices can replace disk drives completely, and removable PCMCIA devices can absorb the functionality of floppy disk drives as well."

60°C pose no problems. PCMCIA flash devices work in altitudes ranging from -200 to +40,000 feet, and in humidity as high as 95% RH. The mean time between failure (MTBF) rating is 250,000 power-on hours versus about 100,000 hours for many disk drives.

■**Performance.** Because flash storage devices are completely solid-state, they eliminate the performance penalties associated with disk spin-up time and rotational latency. When storing operating systems and applications in flash for direct execution by the processor, "instant-on" performance is achieved. PCMCIA flash cards based on high-speed 16 Mbit chips provide read rates of 10 Mbytes/sec, significantly outperforming even the fastest disk drives. Write rates of about .85 Mbytes/sec are proportionally slower, but this factor is overshadowed in typical PC applications where read operations are four times more prevalent than writes.

■**Power consumption.** Flash storage devices consume power only when data is being accessed or written, returning to deep sleep mode the instant the read or write operation is completed. The power spike associated with disk drive motor spin-up is eliminated, and

RAID ENCLOSURES

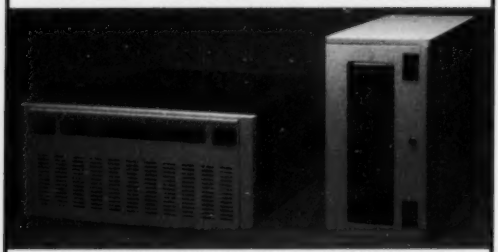

The Sigma SA-H288 and SA-H300 Enclosures for RAID Applications

Before you buy, be sure your enclosure offers the Sigma benefits with your own SCSI, IPI or ESDI fast and fault tolerant subsystem for workstations and minicomputers.

Ask for details on our **NEW SA-H310** Eight-Drive Tower/Rackmount Chassis for the **CMD CRD-5000 RAID CONTROLLER.**

- Hot pluggable 3-1/2" drive support.
- Dual redundant, hot pluggable power supplies.
- Support for internal or external array controller, SCSI string, or software array.
- Complete I/O flexibility, including multiple host support.
- Drive activity indicators.
- External drive and controller ID selection.
- Tabletop, rackmount and tower (even multiple tower) configurations.
- The best pricing available.

Sigma Information Systems

5775 Polaris Avenue, Las Vegas, Nevada 89118
(702) 597-5111 — FAX (702) 597-1017

Call your local Sigma sales office today and find out how you can benefit from the Sigma leadership in RAID enclosures.

CA	(408) 241-3770	CO	(800) 462-5411	NH	(603) 472-9786
	(714) 435-9019	FL	(800) 462-5411	NJ	(908) 499-0014
	(818) 705-7543	IL	(800) 388-9270	NV	(800) 462-5411
		MD	(301) 251-8990	TX	(800) 521-3007

All product and company names are trademarks of their respective companies.

Circle 17 on Reader Inquiry Card

some PCMCIA flash devices require only a 3.3V power source for read and write operations. Some devices can operate with 3.3V or 5V power to allow interchangeability between standard desktop systems and low-power mobile systems. Some PC manufacturers claim twice the system battery life with flash, and some hand-held PC makers are eliminating disk drives completely.

■**Silent, vibration-free operation.** Flash storage devices can be used to design completely silent and vibration free solid-state systems.

The applications of flash mass storage are varied. Flash storage can provide incremental improvements to existing systems by extending performance and battery life. On the other hand, flash storage can take the PC to new environments that were previously too demanding. Today, flash storage systems are deployed in applications including patient monitoring systems, robot control systems, police cars and flight recorders.

An obvious application area for flash mass storage is mobile computing, where the benefits of small size, light weight, ruggedness and long battery life come together. Low-cost, application-specific mobile systems are possible because flash can combine the functions of a hard disk drive, floppy disk drive and shadow RAM.

Flash Storage Options

Flash memory is a building block that can be used to design four kinds of mass storage devices: flash drives, PCMCIA flash cards, peripheral bus add-in boards and resident flash arrays (RFAs) that are placed on the system motherboard (See Table).

Every flash storage device needs to be managed by special software that allows the chips to perform storage functions for the host system. This is also true of mechanical disk drives, which have internal media management controllers.

For integrators and VARs, flash drives are an easy way to design flash mass storage into the system. Here, an on-board CPU and logic circuitry handle all of the software requirements. At the other extreme, customized board-level designs place the entire software burden on the host CPU.

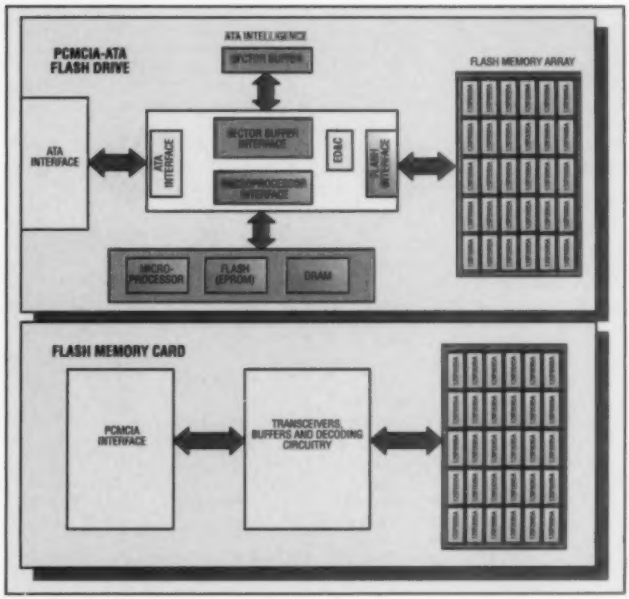

Fig 1 Flash drives, typically packaged with an Integrated Drive Electronics (IDE) interface, resemble a mechanical disk emulator.

Flash drives, typically packaged with an Integrated Drive Electronics (IDE) interface, resemble a mechanical disk emulator. In fact, from the host system's perspective, the flash drive looks exactly like a disk drive, only faster. The circuitry inside the card hides from the system any differences between a disk's rotating platter and flash memory devices (Fig 1).

As the name implies, Integrated Drive Electronics means that the drive contains all the circuitry and software required to run this self-contained unit. In the PCMCIA world, the IDE interface is called PCMCIA-ATA, where ATA stands for AT attachment. Typically, the PCMCIA-ATA electronics occupies most of one side of the PCMCIA card while the flash array populates the other side. The flash array typically provides from 5 to 40 Mbytes of mass storage, depending on the density of the flash chips used.

To shield the host platform from the nuances of flash memory, flash drive IDE logic includes a CPU, an interface controller, a power converter, firmware, a RAM buffer and some miscellaneous components. The microprocessor performs high level IDE command interpretation and handles things like internal data movement and flash media management. The interface controller supports the drive's interface signals and the IDE protocols. It also acts as the mediator between the system bus and the flash drive CPU. Most flash memory devices require a 12V erase/write current, and the drive's internal power converter generates this supply from the 5V system input.

Today, flash drives integrate flash solid-state storage into a system. Similar to the mechanical IDE drive, a flash drive plugs right into the system and is ready to run, provided that a connector is available. Due to the popularity of PCMCIA, the majority of integrated flash drives are designed with the PCMCIA-ATA interface in a Type II card form factor. This allows the card to be plugged into any PCMCIA Type II or III socket. A flash drive designed to comply with the PCMCIA-ATA speci-

fication contains a Card Information Structure (CIS) located in an Attribute Memory Plane. The CIS identifies the PCMCIA card as a flash drive and not a mo-dem or other peripheral. The CIS also describes various drive characteristics including density, manufacturer and power requirements. The CIS promotes dynamic insertion and interchangeability among different PC Cards and different systems.

According to the PCMCIA specification, all compatible drives must power up in the memory mode so that the system can read the CIS and access special configuration registers. A PCMCIA-ATA driver must be installed in the host system to interpret the CIS, set up the configuration, and handle the PCMCIA-ATA protocol. This driver, approximately 8 Kbytes in size, can be installed in a system and loaded via the "DEVICE=" statement in the CONFIG.SYS file. However, in a system with a PCMCIA socket and supporting software, a PCMCIA-ATA driver will typically be included with the BIOS code provided by the BIOS ISV. If you design your own BIOS code, you should integrate an ATA driver that meets the ANSI X3T9.2 ATA specification.

Some drives, such as Intel's iFD010P2SA, can also operate in the true IDE sense by pulling the output enable signal (OE#) low. In the IDE mode of operation, the drive powers up with immediate access to the drive portion itself, not the CIS. This method inhibits interchangeability because without access to the drive's CIS, there is no way for the system to determine whether a modem, flash card, flash drive or other type of PCMCIA card is in the socket.

PCMCIA Flash Memory Cards

There are many PCMCIA flash cards on the market today, and most conform to the PCMCIA Type I format. Like the flash drive, the flash card also contains a flash memory array, CIS and PCMCIA socket interface circuitry. But the similarities stop here. The circuitry is generally in the form of proprietary ASICs that handle functions including address decoding, data bus buffering, status change interrupt configuration and voltage control. The card doesn't contain a CPU, an interface controller, or any of the other major components of the flash drive. Responsibility for flash media management is transferred to the host CPU, and the cost of on-board IDE logic is eliminated.

Because the flash cards interface directly with the host CPU, flash cards can be used to perform eXecute In Place (XIP). In XIP, software applications can be run directly out of the card rather than from DRAM, thereby providing "instant-on" software re-sponsiveness and reducing the cost and board space associated with redundant memory systems. XIP will become increasingly important in hand-held PC applications such as PDAs and Mobile Companions. These systems require new ways of accessing data from the PCMCIA socket and must work across a broad range of CPU architectures.

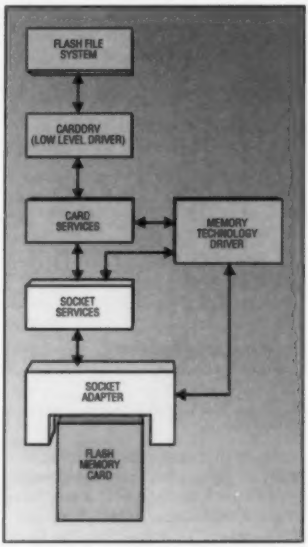

Fig 2 The PCMCIA-ATA architecture is indistinguishable from a rotating platter to the host.

Flash memory cards can be segmented into an XIP partition and a file storage partition. The file storage partition requires a Flash File System (FFS) that allows the system to use the flash card in a disk drive-like manner. FFS software, available from companies including Microsoft and SystemSoft, are special device drivers that optimize the flash memory file storage capability. Whereas disk drives are managed as a collection of sectors, the Microsoft FFS stores files consecutively using linked-list pointers that locate the beginning and end of each file.

Fig 2 displays the PCMCIA software suite used to integrate FFS software with other PCMCIA drivers, the host operating system and hardware resources. These PCMCIA software modules play an essential role in the hot swapping of PCMCIA cards. Hot swapping refers to the system's ability to dynamically reconfigure itself, without removing power, to accommodate I/O cards, memory cards and even flash drives. This ability also allows the system to be auto-reconfigured via a card, a capability sometimes called "plug and play." Card Services monitors the way PC Cards and their applications utilize and interface to a host system. It also handles system resource sharing. Socket Services controls adapter and socket hardware functions. The Memory Technology Driver (MTD) provides the programming and erase algorithms for specific flash memory cards. A low-level driver, CARDDRV, gives API independence to the FFS. After removing a flash memory card and installing a flash drive, the operating system circumvents the FFS and goes directly to the PCMCIA-ATA driver.

An embedded design that does not require card removal and interchange also does not require a complete PCMCIA software suite. The FFS and CARDDRV can be combined to provide the complete solution. In this application, CARDDRV handles the functions of Card Services, Socket Services and MTDs. Although this implies that fewer software pieces are needed, the designer must modify CARDDRV to match the specific hardware requirements of the system and the type of flash memory cards used.

Most BIOS vendors offer a complete PCMCIA software suite. System-Soft also provides Microsoft FFS. You may choose to buy rather than build in order to simplify the task of interfacing system software with flash memory cards.

Analogous to the PCMCIA flash

card with its flash memory array, the flash add-in board provides a relatively simple way of adding solid-state storage to a system. The solution does not provide the flexibility of hot swapping and dynamic system reconfiguration, however. Many companies manufacture add-in boards for the ISA bus, VME bus, or other types of system architecture-specific buses in a variety of flash storage densities. Typically, these systems use mechanical disks as the primary application storage device while the flash memory stores crucial data files and high performance applications, taking advantage of the reliability and speed of solid-state technology. Flash add-in cards are finding application in products like network servers, where network configurations can be downloaded over the network to network management software.

Similar to using a memory card, the host system must provide the software to manage the flash memory array. Although the add-in board doesn't promote the interchangability of memory cards, the same PCMCIA software suite can be used as long as some of the modules are written to accommodate the specifications of the add-in board (Socket Services and MTDs). Socket Services hides the specific hardware configuration of the adapter and MTDs hide the specific flash memory details.

As a cross between a memory card and add-in board, the card drive (sometimes called reader/writer) provides a way to add flash card support to an existing platform. Card drives are available for several bus architectures including ISA, VME and SCSI. Card drives generally employ a cable between the socket and a floppy disk drive bay at the front of the computer, allowing the end user to easily insert memory cards. PCMCIA interface controller chips, such as Cirrus Logic's CL-PD6710, lie at the heart of these add-in boards, allowing a number of commercially available PCMCIA software programs to be used with these drives. In some cases, the card drive vendors supply the software along with their drives.

Resident Flash Arrays

Flash storage may also be designed at the component level. This design is called the Resident Flash Array (RFA) because the flash devices reside on the host motherboard. This approach offers the most flexibility for developing a customized solution for a specific application, but it restricts after-market or integrator add-ins.

When using this approach, the system designer can select the required density, match the bus width to the host CPU, and include various board-level capabilities to optimize the interface. The first task is determining the density required for the specific application. Sometimes, this consideration is impacted by board space constraints. The advanced packaging of flash components makes it possible to achieve very high density board layouts. For example, Intel offers up to 4 Mbytes of storage in a single dual-die TSOP package.

Bus width must be selected carefully. Currently, flash components are available with 8- and 16-bit interfaces. The erase block size for flash memory devices is generally 64 Kbytes. However, when combining several devices in parallel to match a 16- or 32-bit bus, this increases the effective erase block size to 128 or 256 Kbytes.

An important performance consideration is attaching the flash memory to the host CPU. Many OEMs use the local bus approach to achieve the highest performance possible. The local bus is also preferable for XIP applications.

Once again, PCMCIA software may be used. Socket Services can be employed to make transparent the hardware differences between hardwired flash devices and removable memory cards. Other details can be hidden by modifying the MTD to accommodate bus width, control signals and other variables.

"Due to the popularity of PCMCIA, the majority of integrated flash drives are designed with the PCMCIA-ATA interface in a Type II card form factor."

Choosing The Right Flash Storage Solution

Making the right flash mass storage selection is a function of many considerations, with several practical options. Each solution has advantages and disadvantages that should be weighed against the target application requirements and your own design skills. You can count on flash manufacturers to continue to differentiate and proliferate their offerings, giving you a wide spectrum of choices, from plug-and-go to highly optimized solutions.

Markus Levy *is a senior technical marketing engineer for PCMCIA flash memory products at Intel Corp. (Folsom, CA).* **Pedro Vargas** *currently serves as flash market development manager for Intel.*

For more information circle 113

TABLE: COMPARING THE FLASH SOLID-STATE STORAGE ALTERNATIVES

	ATA DRIVES	PCMCIA MEMORY CARDS	ADD-IN CARDS	DEVICE-LEVEL
CONTROL CIRCUITRY	BUILT -IN	HOST-DEPENDENT	HOST-DEPENDENT	HOST-DEPENDENT
FILE SYSTEM SUPPORT	BUILT -IN	HOST-DEPENDENT	HOST-DEPENDENT	HOST-DEPENDENT
SOFTWARE SUPPORT	BIOS VENDORS	BIOS VENDORS	BIOS VENDORS OR CUSTOM	BIOS VENDORS OR CUSTOM
DENSITIES	5-40 MBYTES	1-40 MBYTES	FLEXIBLE	FLEXIBLE
PCMCIA COMPATIBLE	YES	YES	NO	NO
SUPPORTS XIP	NO	YES	DEPENDS ON DESIGN	DEPENDS ON DESIGN
DATA PATH	DEFINED BY PCMCIA	DEFINED BY PCMCIA	FLEXIBLE	FLEXIBLE

RAID

Streaming RAID—A Disk Array Management System For Video Files

Conventional file servers and storage systems are not well-suited to support video over LANs

by Fouad Tobagi, Joseph Pang, Randall Baird and Mark Gang
Starlight Networks

The characteristics of digital video files and traffic differ substantially from those encountered with data applications: (i) video files are much larger than data files, and (ii) video traffic is continuous in nature while data traffic is bursty, with the data rate of a video stream much higher than the mean rate of a data traffic source. Accordingly, conventional file servers and associated storage systems are not well suited to support video services over local area networks; instead, new servers capable of handling the specific characteristics of video files and traffic are needed. This article describes a disk array management system called Streaming RAID which is particularly suitable for a video applications server.

Streaming RAID is a component of StarWorks, digital video networking software based on a client/server model aimed at supporting multiple video users simultaneously. StarWorks consists of server software which turns a standard off-the-shelf hardware platform (such as a PC or a workstation) into a streaming video server, and client software which allows multimedia clients (such as a PC, a Macintosh, or a UNIX workstation) to control video services remotely, and to transmit and receive video stream data over the network. The server software is currently implemented for a PC platform equipped with a 50 MHz i486 microprocessor, an EISA bus, a SCSI-2 bus interface, and multiple magnetic disk drives, and achieves an aggregate stream data rate of 25 Mbps (i.e., supports twenty 1.2 Mbps video users simultaneously). The client software is currently implemented for the PC (DOS 5.0 or higher and Windows 3.1), and the Macintosh (System 7.0 or higher).

Introduction

With data applications, whenever a user makes a file access request to a server, or requests that data be transmitted on a network, the user expects a fast response—fast compared to the time it takes it to place the next request. As a result, the capacity of a server and the overall bandwidth must both be large compared to the average demand placed by a single user. The design of a file server aimed at supporting data applications and the design of a network to support data traffic have thus been based on the principle of bandwidth sharing and statistical time multiplexing. File servers have furthermore taken advantage of the property of locality in file access, and incorporated caching mechanisms well tuned to such access patterns. In all cases, as the overall load on the shared resources increased, the average response time experienced by all users also increased.

A video signal is analog in nature and continuous over time. It is digitized by first sampling it at regular intervals, and then by quantizing each sample. This digitization process results in a data stream which is of

The #1 OEM Choice
Contact CMD for Custom Solutions

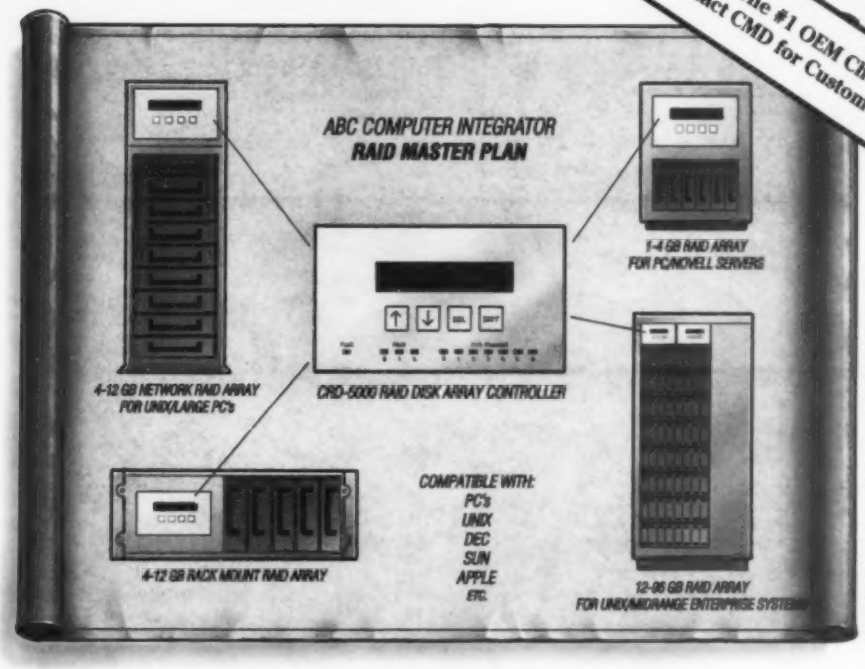

©Copyright 1995 CMD Technology, Inc. All product names are trademarks of their respective companies.

Build the best RAID.

Why settle for less than the best? OEM's, Integrators, and Resellers world-wide are now proclaiming CMD's CRD-5000 RAID 0, 3, and 5 controllers as the best solution for custom-designing their own family of high performance, high availability RAID subsystems.

Whether your target market is low-cost Novell Servers, higher performance UNIX Workstation platforms, or midrange systems such as DEC VAX computers, CMD's CRD-5000 highly scalable architecture fits the bill. Most importantly, the CRD-5000 assures you of driverless operating system compatibility across diverse platforms from Windows NT to HP/UX to AIX to VMS.

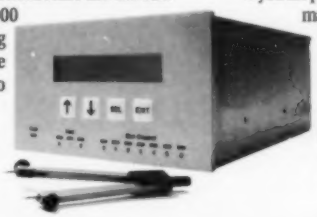

CMD's CRD-5000 incorporates world-class features that allow you to differentiate your RAID offerings from less sophisticated competitive solutions. Here are just a few to whet your appetite: high performance caching using industry standard SIMM's, support for multihosting, multi-ranking, battery-backed write buffers, and a full line of CMD DEC and PC SCSI host adapters that support options like tag queuing.

If you are planning to build your own RAID subsystems, make sure you start with the best foundation. For more information on the CRD-5000 or how to join CMD's RAID Developer program, call today at 1-800-426-3832.

CMD TECHNOLOGY INC. *The Leader in SCSI Interface Technology*

1 Vanderbilt, Irvine, CA 92718 (800)426-3832 (714)454-0800 FAX (714)455-1656

Circle 18 on Reader Inquiry Card

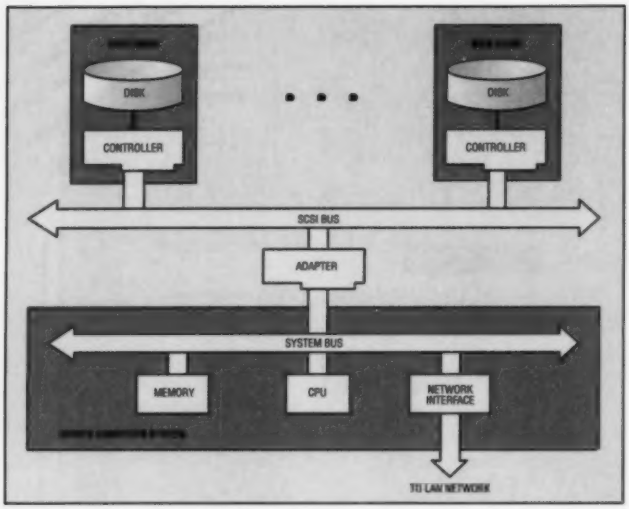

Fig 1 The transfer of video data between the disk array and the network takes place via the system memory.

relatively constant and very high rate; (for example, the digitization of a NTSC video results in data rates in the neighborhood of 100 Mbps and an HDTV signal, 600 Mbps.) However, given that the sampled data exhibits a great deal of redundancy, compression is applied, thus significantly reducing the stream's rate. Depending on the bandwidth of the original analog signal, the sampling rate, the quantization step size, the encoding method, and the desired image quality, the resulting data rate for a digital video signal ranges from 64 Kbps to tens of Mbps. Two important observations can be made:
a) Despite compression, the volume of bits corresponding to a digitized video segment of useful duration is large. A ten minute MPEG video segment requires over 100 Mbytes of storage; ten hours require over 5 Gbytes. Thus, video servers where shared video information is to be stored must have relatively large storage capacity.
b) In order to minimize the client's memory requirement and the latency involved in accessing stored files, the transmission of digital video data between a server and a client is best done in a stream fashion, at the rate at which the data is consumed (in the case of playback) or generated (in the case of recording). In that case, video data must be delivered to the destination on time, and failure to deliver it on time would result in video quality degradation. This has two main implications: (i) from a network's point of view, one requires the availability, on a continuous basis, of a bandwidth at least equal to the signal's data rate; (ii) from a file and storage system point of view, one requires streaming capabilities which guarantee the continuity of each stream being retrieved or stored. In order to support multiple independent video signals, the network must have the necessary aggregate bandwidth as well as means to guarantee the bandwidth required for each video stream, and the file storage system must be of the streaming type and must have a capacity. Conversely, there is a maximum number of video streams of a given data rate that a network and a server can support, and means must exist to prevent additional requests from overloading the system.

The focus in this paper is on the design and operation of a storage system suitable for a video applications server, where the physical storage needs are large, and where, when serving video streams, be it for recording or playback, it is imperative to maintain the continuity of the streams. In the case of playback, data must be retrieved from the storage medium, transmitted over the network, and made available to the decoder no later than the time at which it is needed so as to avoid letting the decoder run dry. In the case of recording, the writing of data on the storage medium must keep up with the rate at which it is getting produced so as to avoid buffer overflow and data loss. Of particular interest is the design and operation of a storage system which:
(i) has the necessary throughput to support simultaneously a target number of users,
(ii) provides simultaneous and independent access to shared files by users, and
(iii) guarantees a start-up latency for new streams within a specified tolerable threshold.

The first step in the design is clearly the choice of a storage medium. Given the above stated objectives, magnetic disk drives emerge as very appropriate for the purpose and are considered in this paper. Indeed, they provide random access read and write capabilities, are readily available, and are of relatively low cost. In order to provide the necessary storage capacity, and/or to achieve an aggregate throughput larger than that achieved with a single drive, multiple drives are used .

Streaming RAID

The transfer of video data between the disk array and the network takes place via the system memory (see Fig 1). Data is transmitted over the network in packets, the size of which is determined by the particular network in use (e.g., 1500 bytes for Ethernet). Given the characteristics of magnetic disk drives and the overhead associated with an I/O transaction data is stored and retrieved on each disk in blocks which we refer to as segments of size larger than that of a packet; as we shall see in the sequel, the block size is a design variable which is determined so as to meet a given desired performance (in terms of throughput and start-up latency).

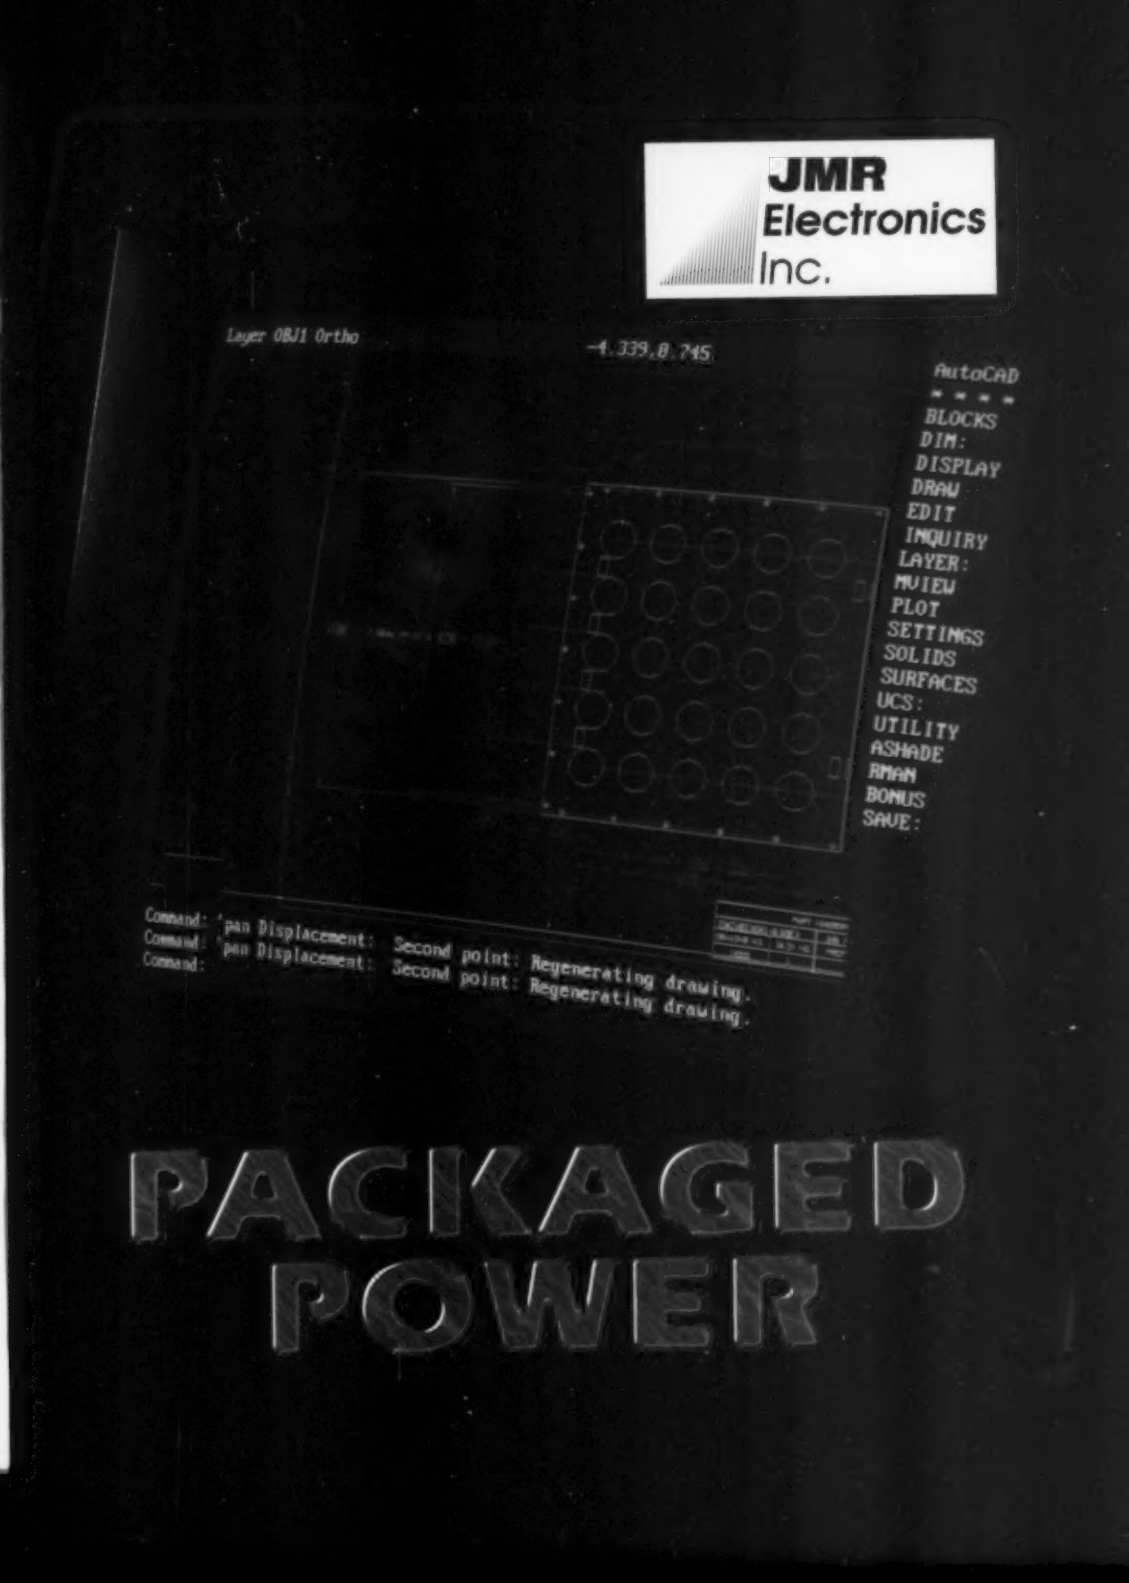
JMR
Electronics
Inc.
Layer OBJ1 Ortho
-4.339,0.745
AutoCAD
BLOCKS
DIM:
DISPLAY
DRAW
EDIT
INQUIRY
LAYER:
MVIEW
PLOT
SETTINGS
SOLIDS
SURFACES
UCS:
UTILITY
ASHADE
RMAN
BONUS
SAVE:
Command: 'pan Displacement: Second point: Regenerating drawing.
Command: 'pan Displacement: Second point: Regenerating drawing.
Command:
PACKAGED
POWER

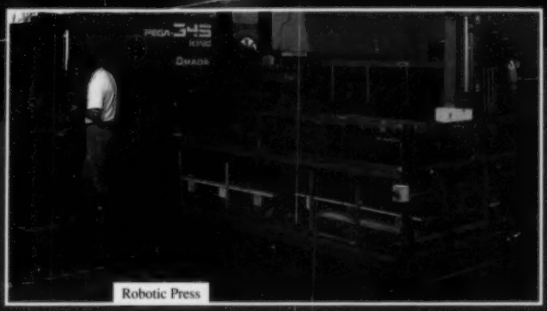

Robotic Press

Robotic Bending Machine

JMR Electronics, Inc. is the world's largest manufacturer of electronic enclosures for computer peripheral devices. As you will see in this catalog, we currently offer more than three hundred different models for almost any peripheral subsystem application. Our focus has always been on supporting our customers with high quality, on time delivery and a high level of customer service. This philosophy has made JMR one of the fastest growing small companies in America.

It is JMR's philosophy that successful products begin with successful engineering. The JMR design team has more than 125 years of combined experience in electrical, mechanical, materials and quality engineering. This skilled group has created an atmosphere where Total Quality Management begins at product inception. Frequent design reviews, coupled with a team engineering approach has increased productivity, quality, reduced waste and shortened manufacturing cycles.

JMR has made a firm commitment to achieving ISO 9000 standards, bringing the highest quality manufacturing to enclosure products worldwide. At the same time we continue to upgrade by automating

Computer Controlled Sheet Metal Stamping

and streamlining our processes to achieve world wide competitive prices. Constant training of all personnel has created an atmosphere where workers can excel. This training supports our goal of ultimately achieving the Malcolm Baldridge Award.

As examples of our upgrades needed to support our dramatic growth over the last years, JMR has made several major additions to our manufacturing centers. The most important upgrade has been our recent move of our main facility from Northridge, CA to a much larger facility in Chatsworth, CA. This new facility houses our metal fabrication shop, paint shop, assembly lines, sales, marketing and administration. In combination with our plastics manufacturing facility (also in Chatsworth) JMR now maintains 125,000 square feet of space.

We have concentrated on automation wherever possible to maximize our ability to produce consistent, high quality enclosure products. Every step in the production of sheet metal parts that requires precision is either computer enhanced (CNC) or fully controlled. JMR's production capacity and quality have now reached world class levels. We can now support production of over 5,000 finished turnkey enclosures

Powder Painting

Mold Shop CNC's

Paint Shop Five Stage Metal Washing

per day. This volume allows us to quickly turnaround hundreds of designs thereby easily surpassing our overseas competitors. Below is a partial description of manufacturing equipment which is currently operating to produce the products shown in this catalog.

Plastic Injection Mold Shop

To support our products which use plastic front bezels, and our line of plastic enclosures, JMR maintains a state of the art plastics facility. Our eight Van Dorn and Aurburg injection mold machines can handle molds requiring up to 30 ounces of material at up to 270 tons of pressure. Plastic parts are assembled using sonic welding and adhesive processes.

To further support our high volume customers, JMR also maintains a full service mold shop where we manufacture custom injection molds faster than any other known mold company. JMR has made molds in one week that would require months to complete at other mold companies. JMR has two Fadal CNC controlled milling machines, EDM machines and a staff of five full time skilled mold designers. With a production force of over twenty full time skilled craftsmen, JMR can produce plastic bezels for almost

Sonic Welding

Plastic Injection Molding Machines

Computer Controlled Injection Molding

any requirement. Because we manufacture the molds in house, JMR can dramatically cut the tooling cost and lead time usually associated with custom plastic parts.

Metal Fabrication Shop

JMR has three AMADA King and one AMADA Queen CNC Turret Punch machines that transform steel and aluminum sheet metal into parts with cutouts for connectors, fasteners, switches, etc. These punches consist of a moving arm that orients the sheet metal, and a turret with a variety of dies that punch the required shaped holes in the metal. The punches are housed in stations on revolving turrets which quickly spin the correct die into position for each punch. For a typical enclosure, each AMADA punch press will punch more than six precision shapes per second.

JMR has recently added a new revolutionary AMADA Robotic Bending brake. This one of a kind brake can be quickly programmed to bend the punched flat metal into almost any three dimensional shape. The robot's mechanical arm retrieves flat metalwork and uses it's internal series of press brakes to bend the metal to exact tolerances. In a high volume manufacturing floor such as JMR's, this piece of

machinery guarantees that the last part produced in a production run is identical to the first without the chance of human error. This reduces the number of inline inspectors and reduces the work load of the Quality assurance inspectors.

JMR's most recent (and most impressive) addition to the manufacturing facility is our new conveyorized automated paint center. This new equipment can apply water or powder based coatings to more than 1500 parts per hour. A variable temperature oven assures an attractive, durable finish. A sophisticated computer control center maintains a consistent texture from the first part to the last. This one million dollar state of the art paint equipment allows us to run two high speed lines to paint powder coat and water based paint at the same time. This new equipment allows JMR to be the first to comply with the tough new 1995 environmental standards of California.

Wire Room

JMR has added a COMAX CNC wiring harness manufacturing station capable of automatically assembling our internal DC power harnesses for our enclosures. This state of the art machine is the only

Quality Control

COMAX Automatic Power Cable Machine

Main Assembly Lines

machine of its type in the state of California. We have also added an automated ribbon cable assembly machine which installs connectors, tests for continuity and hypot before accepting the finished cable assembly, again improving productivity and quality. Completed assemblies produced in the cable room in this manner has been helpful in obtaining and maintaining agency certifications.

JMR Production Staff

JMR has a completely automated MRP system which includes accounting and order entry. This equipment makes JMR Electronics a premier manufacturer of peripheral enclosures, but the equipment is only half of the story. The metal enclosure, power supplies, cable assemblies, switches etc. all come together at our assembly center. The stockroom personnel kit parts via MRP generated work orders and issue parts to the distribution center where all kits are verified. The distribution center in turn issues verified kits to the assembly floor where an automated conveyor belt distributes the kits to the proper assembly team. The team leader at each line manages the assemblers, ensuring that every assembled

Sheet Metal Stock Room

enclosure passes inline inspection as well as QA and Hypot testing. The team leaders are ultimately responsible for quality assurance for the assembled product.

The assembled enclosures are then sent through another inspection by our Quality Assurance department to ensure that proper procedures have been followed in the assembly process and that the finished enclosures match the current engineering documentation. After QA inspection the units are delivered to packaging where they are given a final visual inspection, the covers are installed and the units are packaged for shipment. JMR's quality processes have been certified to meet MIL-Q-45208 standards by Lockheed Corporation and will soon comply fully with ISO-9000 standards.

JMR's staff of over 300 employees working in production, engineering, shipping, purchasing, sales and administration all work together to ensure that our customers achieve a high level of satisfaction with our products and our company. JMR (as of 1993) ships approximately 30,000 enclosures each month with as many as 800 different models and configurations, within budget and to exact customer requirements.

JMR ELECTRONICS, INC. PRODUCT GUIDE–1994

RAID ENCLOSURES

JMR Electronics has the world's most complete line of RAID enclosures. Virtually every situation where RAID devices are used is supported by a JMR enclosure designed to specifically fit that application. Whether the need is for a RAID tower, rackmount, or combination RAID – PC, JMR Electronics has the solution.

Our enclosures are designed to meet the needs of users requiring single or multiple N+1 600 watt power supplies, or single smaller power supplies integrated into each canister (for each device), or combinations of the above.

Currently there are more than a dozen models to choose from and we are constantly adding to our product line. The following is a sample of what JMR is currently offering for RAID drives in the 3½" and 5¼" formats:

RAID 10

The RAID 10 has two N+1 600 watt power supplies and ten hot swappable 5¼" full-height

canisters. Designed to enclose ten drives, this 100 pound heavy-duty enclosure is sturdy, attractive, and mobile. A similarly designed five drive model is also available.

RAID 10/PS

The RAID 10/PS encloses ten 5¼" full-height canisters that contain individual 65 watt power supplies. The data cables

can be configured in various ways to suit specific applications. Space is provided for one or two 5¼" RAID controllers inside the top of the enclosure.

RAID 6N

The RAID 6N has two N+1 600 watt power supplies mounted inside the top of the enclosure. Five 5¼" canisters can be swapped by turning off the canister power switch and pulling the canister forward, removing the data and power connections as the canister is removed.

RAID 5 PCPS

The RAID 5 PCPS model has a 65 watt power supply mounted within each of the five

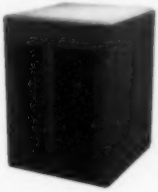

canisters. Each canister has an LED drive ID display. Space is provided for a 5¼" RAID controller and a PC to run as a file server or other configuration.

MINI RAID 10N+1

The MINI RAID 10N+1 encloses ten 3½" drives. It has two N+1 600 watt power supplies mounted inside the top

of the enclosure. Space is provided for two 5¼" RAID controllers inside the top of the enclosure.

MINI RAID 10/TPS

The MINI RAID 10/TPS encloses ten 3½" drives. Each canister has a 60 watt power supply, fan, drive activity LED display, and on/off power supply switch. Space is provided for a 5¼" RAID controller inside the top of the enclosure.

MINI RAID RACK 5PS

The MINI RAID RACK 5PS is a rackmountable RAID enclosure. The enclosure comes complete with five removable 3½" canisters and space is provided for a 5¼" RAID controller. Each canister has its own 60 watt power supply, power on and drive activity LED, quiet fan, and on/off switch.

OTHER MODELS

Currently, JMR offers seven other models of RAID enclosures designed to meet the requirements of different applications. If you don't see a model that suits your particular needs, it may be too new to be listed in this product guide. Additionally, JMR introduces new products every week. Please contact you local JMR sales representative to discuss your specific application requirements.

JMR also offers complete customization capabilities. We pride ourselves in our ability to modify a stock enclosure to suit a customer's needs, as well as our ability to build an enclosure "from the ground up" in order to meet the most demanding specifications. This goes for the rest of JMR's products as well.

We pride ourselves in our ability to manufacture enclosures to the highest quality, competitively priced and in a timely fashion.

JMR Electronics Inc.

TOWER FAMILY OF ENCLOSURES

The Tower Family of enclosures is truly JMR's main line. Currently there are 23 basic models, from which there are virtually hundreds of possible variations.

The Tower Family consists of models for 5¼" full-height and half-height form factors, as well as 3½" standard and 1" height form factors. There are models to house from one to four devices.

Models in this family share some basic characteristics:

All models start with a heavy gauge steel chassis and outer cover. They are available with open front bezels for devices which require access for removable media, or with optional, attractively styled filler panels for devices not requiring access.

Nano Tower

All models come complete with power supplies, cooling fans, SCSI address select switches, and necessary cabling.

Some popular models in the Tower Family include:

The Chico Tower, which

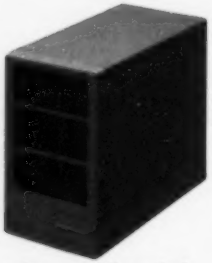

encloses one 3½" standard height storage device. This compact design offers a small footprint for the desktop.

The Cube, which is designed for a single 3½" standard height storage device vertically mounted to provide the smallest footprint possible.

The Micro-Mini Tower, which houses two 3½" devices. One device is a standard height and the other is a miniature (one inch high) device.

Four Plex

The Pico Tower, which encloses two 3½" standard height storage devices, mounted one on top of the other, maintaining a small (less than 6" by 9") foot print.

The Pico-Plex Tower, which is designed to accommodate

four 3½" standard height devices. If required, access is available to the front two of the four drives. The drives are mounted two high and two deep in this enclosure, which measures less than 6" wide by 15" deep.

The Tri Tower, which is designed to house three 3½" standard height devices. These drives are mounted three high, which allows them to fit into a 5½" wide by 10½" deep footprint.

The Pisa Tower, which is designed to enclose four 3½" devices mounted four high, giving this tower the same footprint as the Tri Tower.

The Nano Tower, which is designed for a single, miniature (one inch high) 3½" form factor device.

The Vertical Nano Tower, which is designed to meet the requirements of users wanting to add a miniature (one inch high) 3½" form factor device to their desktop, and maintain the smallest footprint possible.

The Bi-Nano Tower, which is designed to house two miniature (one inch high) 3½" drives in a footprint measuring only 5¼" wide by 8¼" deep.

The Medio Tower, which encloses a single 5¼" half-height device and fits into a footprint only 7½" wide by 10" deep.

The Mesa Tower, which is designed for one 5¼" full-height device or two 5¼" half-height devices.

The Micro Tower, which has a smaller footprint than the Mesa Tower, yet also encloses one 5¼" full-height device or two 5¼" half-height devices.

The Uno Tower, which is designed to meet the needs of users wanting to add a single 5¼" full-height or two 5¼" half-height devices to their desktop. This enclosure mounts the drive(s) vertically, giving the smallest footprint possible.

The Mini Tower, which is designed to house one 5¼" full-height and one half-height or three 5¼" half-height devices in a footprint less than 8" wide by 12" deep.

The Mid Tower, which is designed to enclose two 5¼" full-height devices, or up to four half-height devices (substituting two half-heights for each full-height device).

The Four Plex, which is designed to meet the needs of users needing up to four full-height 5¼" devices or up to eight half-height devices (substituting two half-heights for each full-height device).

PLEX FAMILY OF ENCLOSURES

The Plex Family of enclosures currently consists of nine different models which store from four to eight devices. This family is distinguished from our other enclosures because of the heavy-duty design, large power supplies and device accessibility.

5¼" Penta Plex

These rugged enclosures feature heavy gauge steel chassis and covers. With the exception of the Max Tower, all Plex Family enclosures have removable front panels. The Max Tower has a front door, instead. In either event, the front panels or doors have air vents which allow the powerful fans to circulate air and efficiently cool the enclosed devices.

3½" Octo Plex

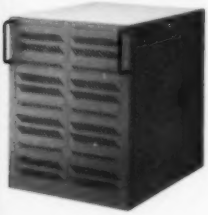

5¼" Hexa Plex

Side doors are standard on all enclosures except the Max Tower. These doors provide easy access to the cabling and rear of each storage device.

The Plex Family consists of four models for 3½" full-height form factors and five models for 5¼" full-height form factor devices. These devices mount on trays, which provide easy removal and replacement of the storage devices.

Front openings can be provided for easy access for removable media devices, and attractive, vented bezels are for devices which do not require front access.

5¼" Octo Plex

The 5¼" models house from four to eight full-height devices. Two mounting trays may be substituted in place of one full-height tray to permit two half-height 5¼" devices to be used instead of one full-height device.

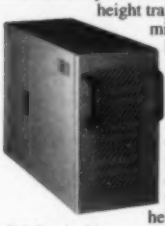

3½" Quadra Plex

All enclosures come complete with power supplies, cooling fans, SCSI address select switches, and interconnecting cabling.

REMOVABLE FAMILY OF ENCLOSURES

5¼" Dual Disc Pac

JMR's Removable Family of enclosures consist of twelve models. These devices are ideal for users concerned about file or software security, and wish to place an entire drive into a safe or other locked environment, or simply wish to transport files the easiest way possible. All enclosures in this family provide shock mounted transportation of storage devices.

5¼" Dual Disc Pac Rackmount

3½" Disc Pac

Duet PS

There are four "Disc Pac" models:

The Disc Pac 3½" model is a rugged 5¼" form factor receiver and canister model which mounts a standard height 3½" drive into any 5¼" half-height mounting slot.

The 5¼" "Disc Pac" models consist of a table top canister and receiver, a dual table top model, and a dual rackmount model. Each model is capable of enclosing 5¼" full-height form factor devices.

In addition to these standard removable enclosures, JMR offers two 2½" enclosures, one is the "Pocket Pac" for Macintosh Mini Single Computers, the other is the "Battery Pac" for general transportability for most computers. The "Battery Pac" is completely self-contained.

Most models in this family include power supplies, quiet cooling fans, SCSI select address switches, and cabling.

These small, state-of-the-art enclosures are ideal for transporting data from computer to computer, and their small size easily allows them to be carried in a briefcase.

5¼" Canister & Receiver (Single Table Top Model)

PLASTIC FAMILY OF ENCLOSURES

The Plastic Family of enclosures is constructed with a heavy gauge steel chassis, enclosed by a molded slide-on cover. Each model of this family comes complete with a power supply, cooling fan, SCSI address switch, and internal data cables.

PL-15

Different models in this family support 2½", 3½", and 5¼" form factors. The 5¼" models support both half-height and

Mini Snap Closed

full-height devices. The 2½" form factor model is designed specifically for Macintosh computers that require a great deal of transportability.

CD ROM FAMILY OF ENCLOSURES

The CD ROM Family of enclosures consist of eight basic models which enclose from one to eight 5¼" half-height CD ROMs.

Models include:

The Mac Zero Low Profile Open, which has a heavy gauge metal cover and chassis, and is designed for one half-height 5¼" CD ROM.

Info Tower 7

The CD40 through CD80 series, which enclose from four to eight half-height devices.

The Info Tower 5 and Info Tower 7, which have attractive, lockable front doors for security.

All models in this family are equipped with power supplies, cooling fans, SCSI address switches, and internal data cables.

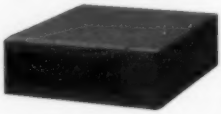

HD2W

ZERO FOOTPRINT FAMILY OF ENCLOSURES

Designed for either 3½" standard height, or 5¼" half-height or full-height devices, the Zero Footprint Family models are an attractive addition to any desktop. Like most of JMR's products, these models are completely turnkey, in that they come with power supplies, cooling fans, SCSI address switches, and data cables.

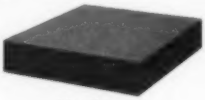

Mac Zero
Low Profile Closed

WORKSTATION ENCLOSURES

All models in the Workstation product line are constructed from heavy gauge steel chassis and coverings. There are models designed for 3½" standard-height or 5¼" half-height or full-height form factors.

Open or closed front bezels are available to suit the user's requirements. All models are complete with power supplies, cooling fans, SCSI address switches, and internal data cables.

Metal 5¼" Filler Panel

RACKMOUNT ENCLOSURES

JMR has thirteen basic models of Rackmount enclosures (not including Removable Drive Rackmounts or Industrial PC Rackmounts). These models are

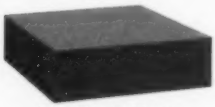

Solar Plex 1H Closed

designed specifically for two or more half-height or full-height devices; most models are available with either open or closed front bezels. All models are extremely durable and rugged in design and are complete with rugged slide rails, power supplies, cooling fans, SCSI address switches, and data cables.

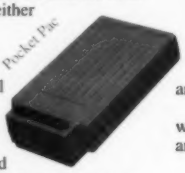

Pocket Pac

RM 300

OTHER ENCLOSURES AND ACCESSORIES

JMR has a complete line of miscellaneous enclosures designed for various needs from

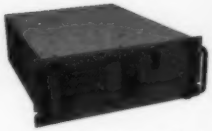

RM 7000 Industrial PC

floppy drives to removable media storage devices. Additionally, cables, terminators, mounting brackets and filler panels are available to meet just about any possible requirement.

1SV5

If you don't see the enclosure you need, please contact your JMR representative for further information on our products

Centronics Cable

and accessories.

In all probability, JMR has what you need, or will design and build it for you.

Highest Quality
Since 1982
American Made
On Time Delivery
Priced Right
Over 1,000,000 Sold

JMR Electronics Inc.

F O R M O R E I N F O R M A T I O N

CONTACT YOUR JMR REPRESENTATIVE OR DISTRIBUTOR TODAY

Corporate Headquarters.....20400 Plummer Street, Chatsworth, CA 91311 • (818) 993-4801 • Fax (818) 993-9173
Northern California.....30 Railroad Avenue, Suite 2, Point Richmond, CA 94801 • (510) 215-1800 • Fax (510) 215-1820
East Coast Office.....1 Harbour Place, Suite 105, North Myrtle Beach, SC 29582 • (803) 280-0405 • Fax (803) 280-3082
European Headquarters.....9 Avenue Coghen, 1180 Brussels, Belgium • 322-344-9542 • Fax 322-344-5438
United Kingdom Office.....Unit 9, Chiltern Enterprise Centre, Station Road, Theale, Berks, RG7 4AA • 0734-304985 • Fax 0734-304546

JMR Electronics Inc.

OVER 300
STANDARD ENCLOSURE MODELS

Steel, Plastic, Steel & Plastic Combinations
Turnkey or Completely Customized

YOUR BEST SOURCE
FOR PERIPHERAL ENCLOSURES

CALL YOUR SALES REPRESENTATIVE TODAY!

East Coast .....(803) 280-0405 – Fax (803) 280-3082
North Central Western .....(818) 993-4801 – Fax (818) 993-9173
South Central Western .....(818) 993-4801 – Fax (818) 993-9173
Northern California.....(510) 215-1800 – Fax (510) 215-1820
Southern California.....(818) 993-4801 – Fax (818) 993-9173
Latin America .....(818) 993-4801 – Fax (818) 993-9173
Europe .....0734-304985 – Fax 0734-304546
United Kingdom.....322-344-9542 – Fax 322-344-5438

Fill out the card to receive more information

DON'T DELAY MAIL TODAY

Please send information on the following enclosures:

- ☐ RAID
- ☐ CD ROM
- ☐ Plastics
- ☐ Removable Drives
- ☐ Music Industry Frames and Rackmounts
- ☐ Towers
- ☐ Rackmounts
- ☐ Sun Workstation Storage
- ☐ Macintosh Zero Footprint
- ☐ Industrial
- ☐ Medical
- ☐ Instruments

Name _____
Title _____
Company _____
Address _____
City _____ State _____ Zip _____
Phone # (_____) _____
Fax Phone # (_____) _____

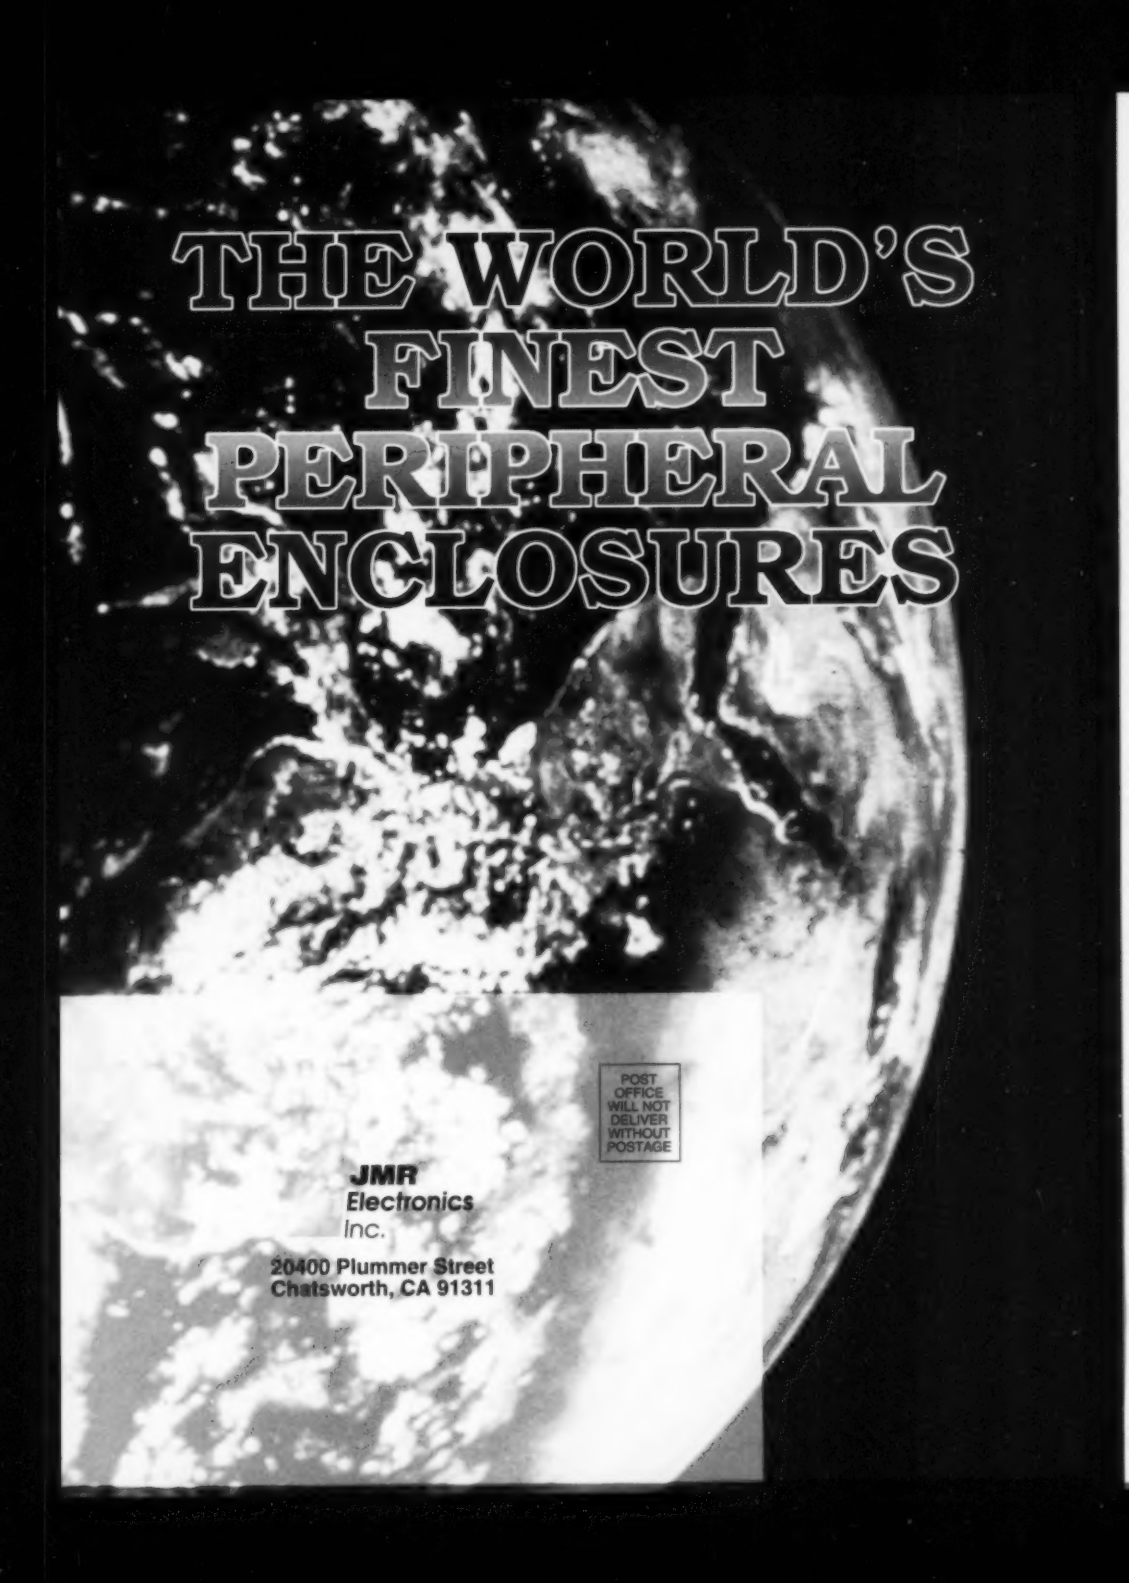
THE WORLD'S
FINEST
PERIPHERAL
ENCLOSURES
POST
OFFICE
WILL NOT
DELIVER
WITHOUT
POSTAGE
JMR
Electronics
Inc.
20400 Plummer Street
Chatsworth, CA 91311

Given that the focus in this paper is on the design of the storage management system, we ignore any network effect (delay or loss), and thus assume that the video data is produced (in case of recording) or consumed (in case of play-back) locally. We also consider in this paper that the data rate for a stream is constant and is the same for all streams; we denote the stream's data rate by W. For ease of implementation, we also consider the system to operate with segments of a fixed size , and denote the segment size by S. The primary interest is to devise a disk access algorithm which maximizes the number of streams, N_s , that can be supported simultaneously, given the available memory size and a desired start-up latency. Conversely, we are interested in determining the required resources (namely, the number of disks and the size of system memory) in order to support a given number of streams. (Without loss of generality, we consider only the case of play-back.)

Consider first the case of a single disk drive. The simple fact that there has to be one I/O transaction per stream every S/W seconds implies that, when the system is serving the maximum number of streams, it has to operate in a cyclic (periodic) fashion, such that the length of a cycle is equal to S/W seconds, and in each cycle there is one and only one I/O transaction per stream. Thus, the number of streams which can be supported simultaneously (N_s) is determined by the number of I/O operations which can be performed during a cycle, and is maximized by minimizing the time it takes to perform these I/O operations. Since the time it takes to perform a certain number of I/O's is a direct function of the segment size S, the physical location on the disk of the various segments that are needed, and the I/O overhead incurred in each I/O transaction, N_s is also dependent on these factors.

In order to maximize N_s, it is important to minimize the effect of I/O overhead. Little can be done about the rotational latencies since the access pattern to video data from the users is unpredictable. The overhead due to the seek times, on the other hand, can be decreased by pooling together several I/O transactions, and by performing them in an order which corresponds to their locations on the disk; (this method is well known and is commonly referred to as the elevator scheme). To apply this method to our situation, we operate the system in a gated fashion, whereby the consumption of segments for all streams are time-aligned, and whereby the segments to be consumed in a cycle are fetched in the preceding cycle. (Fig 2.)In order to support more streams than possible with a single disk, multiple disks are used. Let N_d denote the number of disks. In order to achieve the same good performance per disk as for the single disk case, we operate each disk in the gated fashion described above. The question then arises as to how to allocate the video files to the disks.

To remove any restriction on the users as far as file access is concerned, and achieve with the N_d disks the N_d-fold increase in performance independently of the file access pattern, it is important that the load on all disks be the same at all times (i.e., it is important to guarantee that, in each cycle, an equal number of I/O's are requested from each disk, no matter what the file access pattern is.) A simple way to accomplish this is to consider the N_d disks to form an array, and to stripe the video files across all disks; that is, we consider the segments in a file to be grouped into data stripes of N_d consecutive segments each, to consider segment bins in corresponding locations on all disks to constitute bin stripes of N_d bins each, and to allocate data stripes to bin stripes the same way segments of data are allocated to segment bins in a single disk. The length of a cycle becomes now N_dS/W seconds, and in each cycle N_d segments are fetched for each stream, one segment from each disk, (thus guaranteeing that all disks are accessed equally) (Fig 3). The maximum number of streams that can be supported by a disk array of N_d disks is the number of I/O transactions that can be supported by each disk during a cycle of length N_dS/W seconds. Similarly to the case of a single disk, the determination of this number may only be determined subject to a bound ε on the probability that any one stream experiences a discontinuity.

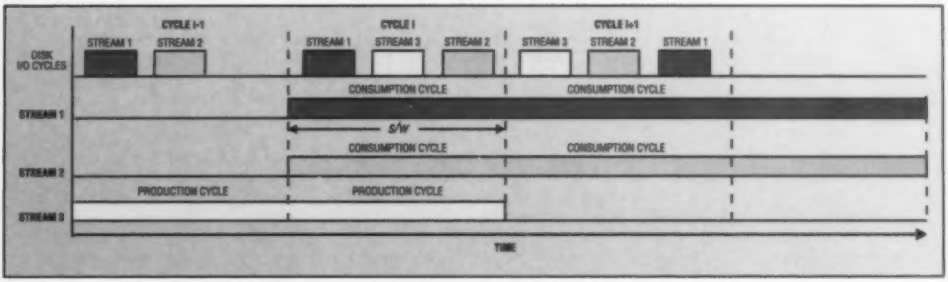

Fig 2 The consumption of segments for all streams is time-aligned; the segments to be consumed in a cycle are fetched in the preceding cycle.

Memory Requirement

The data fetched from the disks is stored in memory until it is consumed. We examine in this section the memory requirement M that is needed for the operation described

above, under the assumption that segments are consumed continuously at the rate of one segment every S/W seconds. M is function of the method used to allocate the memory to the various streams; several methods are considered here, leading to different memory requirements.

The first method is referred to as Complete Partitioning whereby the memory is partitioned into separate regions, one for each stream. Given the gated operation described above (in which the consumption of stripes by all streams are time-aligned), every stream has N_d segments that are ready to be consumed at the beginning of each cycle and N_d requests that are outstanding. Since the order in which the I/O's are performed is not known a priori, and it is possible for all segments requested from the disks for a given stream to be transferred at the same time, N_d buffers have to be allocated to each stream for the requests placed at the beginning of the I/O cycle; thus $2N_d$ segment buffers are needed per stream, leading to a total memory size requirement given by $M=2N_sN_dS$

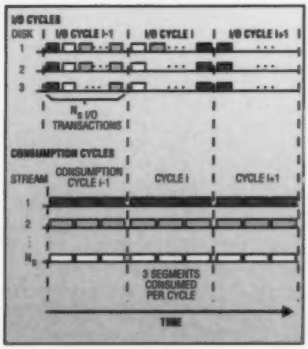

Fig 3 In each cycle N_d segments are fetched for each stream, one segment from each disk, guaranteeing that all disks are accessed equally.

The second method is referred to as Memory Sharing; it leads to a decrease in the required number of buffers by allowing the buffers to be shared among the streams. Further decrease in M may be achieved at the expense of added complexity (and possibly some performance degradation). The idea is to divide the streams into several groups, say G, and divide the I/O cycle into G subcycles, one for each group. Service is still gated, and within each subcycle, the corresponding I/O transactions are sorted so as to minimize the seek overhead. The consumption cycles for the various groups, however, are offset with respect to each other by a subcycle so that the maximum memory requirement for a group, which corresponds to the beginning of the consumption cycle for that group, does not coincide with that of other groups, thus decreasing the total memory requirement. The possibility of degradation in performance is due to the fact that, with a smaller number of I/O transactions per subcycle, the savings in overhead due to sorting the I/O requests is not as high as with the single group case.

It is clear from the above discus-

Trimm Offers an Oasas™ in the RAID Integration Desert.

Why wander around the great Integration Desert when you can get all the power and ease of Integra OASAS RAID Software ready to use with [illegible] RAID enclosure. With this software, simply connect your favorite drives and integration is complete in minutes.

Integra OASAS RAID software has been specially enhanced to take advantage of the unique features of Trimm RAID enclosures including monitor and control fan speed, therm[illegible] sensing and reporting, alarms for component failure–plus all the standard features of OASAS software and Trimm enclosures.

One enclosure. One set of RAID software. O[illegible] stop shopping for integrators. *It's no mirage–it[illegible] the latest in innovative integrator tools from Trimm*

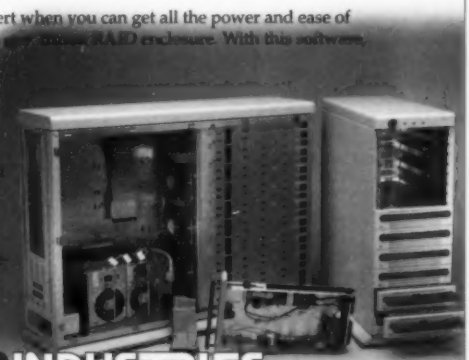

INTEGRA TECHNOLOGIES INC.

TRIMM INDUSTRIES
A Subsidiary of ROBROY INDUSTRIES, INC.
800-423-2024 FAX: 818-503-0438

Circle 19 on Reader Inquiry Card

sion that M is function of the number of streams supported, the number of disks in the array, and the segment size; the larger any one of these parameters is, the larger M is.

Start-Up Latency

Start-up latency is the time elapsed since a request to initiate a new stream is made until the stream may begin to consume its data. Two factors contribute to this latency. First, a stream requires a certain amount of data to be made available before it is able to begin data consumption; (we refer to this data as the start-up data.) Secondly, once the start-up data has been retrieved, a stream has to wait for the beginning of its next cycle before it can begin consuming its start-up data; (recall that in the gated mode of operation, all streams have their consumption cycles time-aligned).

The start-up latency incurred by a new stream is dependent on many factors; namely,

- the amount of start-up data needed for that stream;
- the number of streams that are active at the time of request, as this directly affects the available disk access capacity at the time of arrival;
- the number of streams which are in the start-up phase at the time of request, and the amount of residual start-up data for each,
- the number of new streams which arrive during the start-up period of the tagged stream, and possibly the arrival process according to which these streams arrive,
- the service discipline according to which the new streams are served,
- the scheduling discipline according to which active stream I/O's and new stream I/O's are scheduled in a cycle.

Clearly, the assessment of the start-up latency is a complex task. For the purpose of this paper, suffice it to say that one main indicative factor of latency is the cycle length, since the longer the cycle length is, the longer a stream will have to wait to get its start-up data and start playback.

Numerical Results

In Fig 4, we plot the maximum number of streams as a function of the segment size for various number of disks in the array. We note that the number of streams increases with the segment size, since as the segment size increases, the overhead per I/O transaction gets amortized over a larger transfer time. However the increase levels off as the segment size gets large and reaches asymptotically the disk throughput limitation.

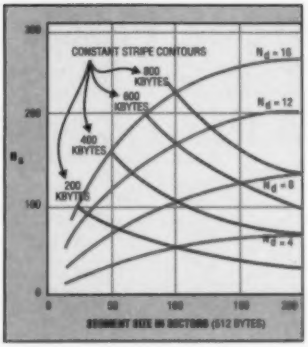

Fig 4 The maximum number of streams can be plotted as a function of the segment size for various number of disks in the array.

Keeping the segment size fixed, the number of streams that can be supported increases with the number of disks in the array. The increase is slightly more than linear with N_d owing to the fact that with the larger number of I/O transactions per cycle, a higher degree of statistical averaging is attained. We note, however, that the increase in the number of streams comes at the expense of an increase in both the cycle length and the buffer requirement.

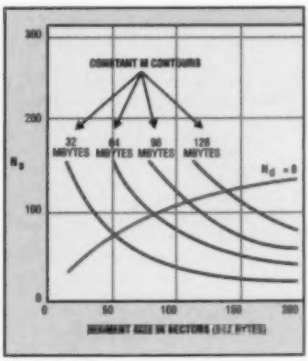

Fig 5 As far as the number of streams is concerned, when S is small it is the disk performance that is the limiting factor; as S increased, memory size becomes the limitation.

It is possible to maintain the cycle length the same if the segment size is decreased so as to maintain the stripe size N_dS constant. We plot in Fig 4 the constant stripe contours; moving along one such contour, one sees that the increase in the number of streams for a given increase in the number of disks is not as large as the increase in the number of streams maintaining the segment size the same. Thus there is a penalty for having to maintain the cycle length the same.

We now address the memory requirement M. Since M depends on N_d, M must be plotted for each value of N_d separately. In Fig 5 we plot N_S and constant M contours versus S for the case N_d=8. We note that as far as the number of streams is concerned, when S is small, it is the disk performance that is the limiting factor, and as S increases, it is the memory size which becomes the limiting factor. Given N_d and a memory size, the maximum number of streams that can be supported by the system is given when the optimum segment size is selected, at the intersection of the N_s curve and the constant M contour.■

Fouad Tobagi *is chief technical officer at Starlight Networks Inc. (Mountain View, CA) and professor of electrical engineering at Stanford University.* **Joseph Pang** *is a member of the research and advanced development group.* **Randall Baird** *is a principal architect, and* **Mark Gang** *is vice president of engineering at Starlight.*

For more information circle 115.

PCMCIA

Block Flash File Systems Deliver Mass Storage For Mobile Computing

By emulating a disk at the sector level, Block FFS makes flash memory an ideal solution for mobile mass storage

by Amir Friedman
M-Systems Flash Disk Pioneers

Flash memory's advantages of high capacity, non-volatility without battery, and low power operation go hand-in-hand with its inherent disadvantages of relatively high price and device write limitations. But, there are solutions available.

Is It A Disk Or Not?

A Flash Disk consists of hardware and software that emulates a rotating magnetic disk using flash memory solid-state technology. For full compatibility with software applications and utilities such as compression and disk utilities, disk emulation must be done at the lowest level possible, the sector level.

PCMCIA flash memory Cards are common, but it is only when they are integrated with the necessary Block Flash File System (Block FFS) that they become Flash Disks. There are several approaches to Block FFS; the TrueFFS system developed by M-Systems illustrates one way to adapt flash memory to the needs of the removable storage market.

Market Acceptance

Two things must happen to ensure Flash technology's acceptance in the PCMCIA market. The first is lower prices, and the second is plug-and-play capability across PCMCIA platforms.

Price reduction is a Catch-22 dilemma: it requires high volume (i.e. market acceptance), but market acceptance requires low price. Another leverage point is needed, which is offered by embedded systems, a large market where flash memory is gaining ground as an EPROM, EEPROM or even SRAM replacement. Using the same type of flash memory technology that is targeted at the mature embedded market the PCMCIA market can achieve the high volume needed to overcome the Catch-22.

Price also depends on the number of sources producing flash memory, so PCMCIA Flash Disks should be able to incorporate flash memory devices from as many manufacturers as possible. For this to become a reality, any Block FFS software must support all types of flash memory architectures and manufacturers.

The second issue involved in market acceptance is plug-and-play capability. Standardization plays an important role here, and PCMCIA is the organization responsible for standards in this environment. One of the more important issues when standardizing a storage solution is the data format placed on the medium. M-Systems has proposed the TrueFFS data format to the PCMCIA committee for standardization.

This underlying data format is compatible with any flash memory technology available today, and will be compatible with any future developments of flash memory technology. This is a very important characteristic of any standard, for no flash memory chip manufacturer should gain an unfair advantage over another through a standardizing process.

Block Flash File System

Block Flash File System is a term used to describe flash memory management software that performs the task of fully emulating a random-access read/write block storage device to the host system.

A file system is a method of orga-

DISCOVER

PCMCIA

100% ExCA™/ "QuickSwap" Compatible

On Your Desktop

The SCM SwapBox—

You've probably already heard a lot about PCMCIA but did you know that you can take advantage of this new and exciting technology on your very own desktop PC? Introducing SCM's SwapBox family of PCMCIA desktop card drives—the industry's first PCMCIA Rev. 2.X card drives made with the desktop computer user in mind.

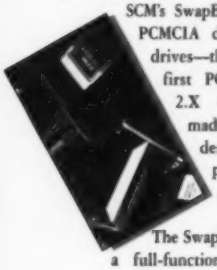

The SwapBox provides a full-function peripheral socket which will allow you to plug in and use NOT ONLY standard memory cards, BUT ALSO fax/modems, LAN adapters, and many many more PCMCIA cards. For many, the SwapBox offers the ultimate "bridge" between your desktop PC and any PCMCIA equipped portable computer, allowing you to share the wide variety of function cards available. You not only save on the cost of peripherals, but you secure your current desktop purchase by linking to the new world of PCMCIA.

With SwapBox's user-friendly software and high-performance hardware, the desktop computer user can take full advantage of all the intrinsic features and functions that PCMCIA offers: Small form factor peripherals *Ability to easily share peripherals between platforms *Removable data storage providing ideal security *Portable technology that's rugged and reliable *Access to hundreds of multi-function PC Card solutions.

SCM offers an extensive line of both internal and external solutions to accommodate many different desktop environments. Give your SCM sales representative a call today to learn more about the advantages of PCMCIA and the many options available to you. Call Today at (408) 395-9292.

IN THE U.S. AND CANADA
SCM Microsystems Inc.
985 University Ave., Suite 7
Los Gatos, CA 95030
Phone (408) 395-9292 • Fax (408) 395-8782

IN EUROPE
SCM Microsystems Gmbh
Fraunhoferstraße 11a
82152 Martinsried, Germany
Phone +4989 859-8702 • Fax +4989 859-5806

SCM MICROSYSTEMS

Circle 64 on Reader Inquiry Card

©1994 SCM Microsystems, Inc. SwapBox is a trademark of SCM Microsystems, Inc. All other trademarks are the property of their respective holders.

nizing data on mass-storage devices into files and directories, and of then accessing such data. It is usually an integral part of an operating system, independent of storage device type. It will work on fixed disks, floppy disks, tapes, optical disks, etc. The file system makes one assumption on the device it is working with — that it is a random-access block device.

Why is it that the standard file system of an operating system cannot be used to manage flash memory-based media? Simply because flash memory is not a random-access device: it can be randomly read, but not written. In addition, flash memory erase block sizes often do not match the sector size associated with disks.

The first issue is making the Flash Disk appear as a contiguous array of fixed-size sectors that can be accessed by the standard file system of the host computer. Since the physical flash devices cannot be randomly write-accessed, the Block Flash File System must create a virtual array of fixed-size sectors that is presented to the host computer as a random-access array of sectors. The host can freely read and write to any sector it chooses without any limitations.

After creating this virtual array of sectors, the Block FFS must keep track of the association between the virtual sectors and the physical sectors. This association is called the virtual to physical mapping (Fig 1). A sector number in the virtual array does not necessarily coincide with the same sector number in the physical array; the Block FFS is responsible for keeping track of which virtual sector is located in a specific physical sector.

Garbage Collection

This mapping scheme is necessary because a given location in flash memory must be erased before it is rewritten. Because of the mismatch between flash memory and disk sector sizes, a sector cannot be erased and rewritten without affecting other sectors around it, so the Flash File System can only update a sector by writing it elsewhere, making a note that the virtual sector's location has changed. Therefore, at any time, flash memory space can be divided

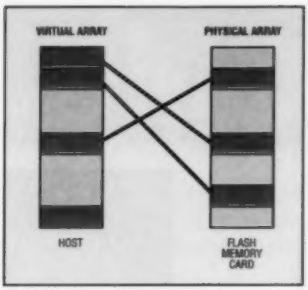

Fig 1 A Block FFS system must map the association between the virtual sectors and the physical sectors, since a sector number in the virtual array does not necessarily coincide with the same sector number in the physical array.

into three types:

- Space where valid data is stored
- Free space
- Unusable or "garbage" space

Free space is depleted as file system operations (copy, delete, rename, etc.) are performed, so eventually garbage collection—converting garbage space into free space—must be performed.

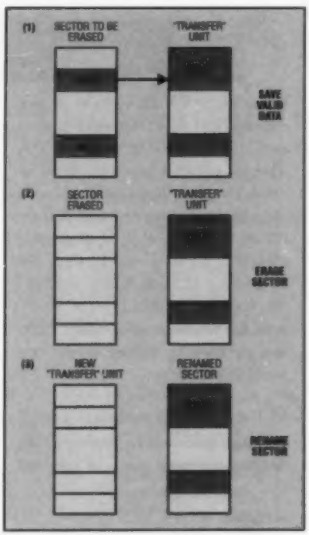

Fig 2 Since free space is depleted as file system operations are performed, garbage collection must be used to covert unusable space into free space.

The Block FFS always has a spare area set aside for the garbage collection process. This area is called a "transfer" unit.

The garbage collection process follows the following steps (Fig 2):

1. Copy the valid data from the sector to be erased to the transfer unit.
2. Erase the sector.
3. Rename the newly erased sector as the new transfer unit.

The net effect of garbage collection is to convert all garbage space on the selected sector into free space. Now the Block FFS has new free space it can work with and it can write more data, until it again runs out of free space. At that point, the garbage collection cycle begins again.

File-Level Flash File System

Another solution is a file-level Flash File System. This emulates a disk at the file level, unlike the Block FFS which emulates at the sector or block level. While the Block FFS works in conjunction with the operating system's file system, the file-level FFS replaces it. Files are stored on the flash media in a linked-list format, whereby each record points to the next record in the file. This type of FFS is implemented in Microsoft's FFS-2 product.

ATA vs. Block FFS

There are two types of PCMCIA Flash Disks, ones based on ATA (AT-Attachment, or IDE) interface, and Linear Flash Cards used in conjunction with a Block FFS. The ATA Flash is a hardware implementation of the function performed by the Block FFS in software. An operating system views the ATA Flash similar to an IDE disk drive.

Fig 3 shows the hardware and software architecture associated with each approach. Functionally, both approaches emulate a disk drive at the sector or block level, and both hardware solutions reside similarly under their appropriate device driver.

The Table shows some of the differences between the two approaches.

Both solutions are viable PCMCIA storage solutions, but the Block FFS-based Flash Disk has the advantage of operating at times not as a disk drive, but as a memory device. This enables it to execute programs

Speed.

SINCE 1988, CPU performance gains have out paced those of conventional rotating disks more than ten to one. So, unless you're already using MegaRam solid-state disks (SSD), your storage strategy isn't keeping up with the times.

MegaRam SSD offers the performance you need to take advantage of today's fastest processors.

- 1/8th Millisecond Access
- Data Transfer Rates Up To 20 MBytes Per Second
- Multi-Porting
- Complete Disk and Battery Backup Options
- Redundant Power Supplies

Since 1978 we've been busy developing the most full-featured, cost-effective systems available. Today, there are thousands of MegaRam systems installed worldwide. We offer solid-state disk configurations for almost all CPU's.

Now, while we didn't write the book on speed, we do have an enlightening brochure on the fastest solid-state disks available and how they fit into your data storage strategy. Simply call us toll-free and we'll rush you a copy.

800-451-0666

IMPERIAL

Imperial Technology, Inc.

2305 Utah Avenue • El Segundo, CA 90245
Telephone: (310) 536-0018 • FAX: (310) 536-0124

Circle 21 on Reader Inquiry Card

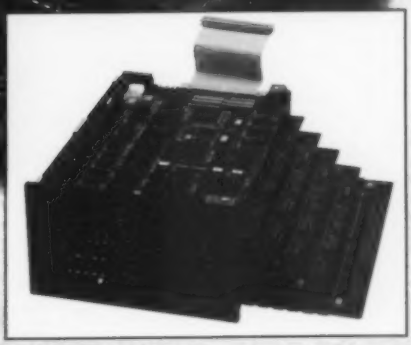

MegaRam-SCSI Component. Start with as little as 16MB. When you need additional capacity, simply slide in another memory module; 16MB, 32MB or 64MB each. Up to 320MB of solid-state storage fits into in a 5-1/4" form factor.

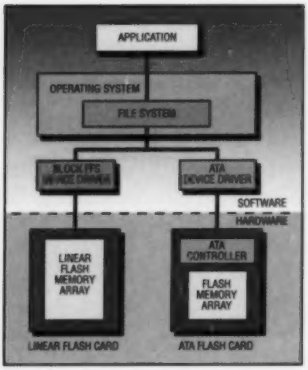

Fig 3 The hardware and software architectures of both ATA and Linear Flash Card-based flash disks emulate a disk drive at the sector or block level and reside similarly under their device driver.

directly off the card without having to initially load them into the computer's memory. TrueFFS and XIP partitions can coexist on the same card. The added cost of the ATA Flash Disk at the low capacity range is due to the microcontroller required to emulate the ATA interface. At the high capacities, the cost of this controller is offset by the cost of the Flash devices.

TABLE: ATA VS. BLOCK FFS BASED FLASH

CHARACTERISTIC	BLOCK FFS-BASED LINEAR FLASH DISK	ATA FLASH DISK
INTERFACE TO COMPUTER	MEMORY	I/O
HOST SOFTWARE SUPPORT	BLOCK FFS DEVICE DRIVER (TrueFFS)	ATA DRIVER
OPERATE LIKE DISKS	YES (WITH TrueFFS)	YES (WITH ATA DRIVER)
EXECUTE IN PLACE (XIP) COMPATIBLE	YES	NO
PERFORMANCE	HIGH DUE TO MEMORY INTERFACE	LOWER DUE TO I/O INTERFACE
COST	LOWER IN LOW CAPACITY CARDS, SIMILAR AT LARGE CAPACITY	HIGHER IN LOW CAPACITY CARDS, SIMILAR AT LARGE CAPACITY
MANUFACTURERS	ABOUT 20	ABOUT 5

MobileMax Flash Disk

Maxtor's MobileMax Flash Disk illustrates the advantages of Block FFS-based Flash Disk technology. It incorporates M-Systems' TrueFFS software and hardware in the form of a PCMCIA Flash Card with plug-and-play capability on many PCMCIA platforms. The MobileMax Flash Disk will operate on systems that are either PCMCIA Card Services or Socket Services software compatible. In the event neither Card Services nor Socket Services are installed in the computer, the Flash Disk will automatically install the correct Socket Services software needed.

The MobileMax Flash Disk has built-in enhancements to target the mobile hand-held computer market. Even though flash devices requiring 12 volts are used in the product, an on-board DC to DC converter is automatically enabled when the Flash Disk is plugged into a 5 volt-only computer. This enables the Mobile-Max Flash Disk to perform read as well as write operations on 5 volt-only computers, where a standard PCMCIA Flash Card using 12 volt flash devices is capable of read only.

Enhanced Address Generation for those hand-helds having a limited addressing capability is also built into the MobileMax Flash Disk. For example, if an HP-95LX is limited to only 2 Mbytes of address space (i.e. SRAM cards), paging techniques are used in the MobileMax product to enable access to the maximum allowable PCMCIA storage capacity (64 Mbytes).

Another function which greatly enhances performance and power savings, is the capability to interleave as little, or as many flash devices as required. Hardware interleaving, in this case, is a method of accessing several flash devices at once, thereby increasing data throughput. Interleaving presents a trade-off between performance and power savings. If quick access to data is important, the interleaving can be set to access several flash devices at once (i.e. effectively increasing data bus width). On the other hand, not all hand-held systems running on batteries can deliver the amount of current needed for this. Those systems can work at the lowest interleaving setting: a single Flash device accessed at any one time. ■

Amir Friedman *is vice president for U.S. Operations for M-Systems Flash Disk Pioneers (Melville, NY). He is also a director and founder of the company which is based in Tel Aviv, Israel. He holds a B.Sc. in Electrical and Computer Engineering from the University of Wisconsin-Madison, and an MBA from Tel-Aviv University.*

For more information circle 107.

Beyond RAID

Save a fortune in computer and manpower time. Introducing Storage Computer Corp's RAID 7™ data storage subsystem.

RAID 7 is a quantum leap beyond lower level RAID implementations that provides virtual solid state disk performance using commodity SCSI disk drives.

RAID 7 shatters I/O bottlenecks with storage power that can double, triple or even quadruple the application performance of your existing CPUs. RAID 7 provides scalable fault tolerance, which means you'll never worry about losing data. And RAID 7's open standards based architecture lets you expand capacity at the lowest cost per megabyte.

to Fault Tolerant

Indeed, RAID 7 is a remarkable storage breakthrough that lets you share your disk investments across different platforms from DEC, IBM, SUN, HP, NCR and others.

Virtual Solid State Disk Storage

Now when you grow or migrate your storage investment, RAID 7 moves with you.

For more information on how you can get the benefit of solid state performance at 1/25th the cost, call us today at 603-880-3005 or fax 603-889-7232.

STORAGE COMPUTER CORPORATION

Authorized VARS and Integrators

▪ DISK DRIVES ▪ TAPE DRIVES ▪ OPTICAL DRIVES ▪ WORKSTATIONS ▪ CONTROLLERS ▪ CABLES

▪ DESKTOP COMPUTERS & NOTEBOOKS ▪ MONITORS/(X)TERMINALS ▪ POWER PROTECTION ▪ MEDIA ▪ SCANNERS ▪ UNIX TRAINING

▪ MODEMS/COMMUNICATION ▪ SOFTWARE ▪ NETWORKING CARDS ▪ GRAPHIC CARDS ▪ ENCLOSURES ▪ ETHERNET CARDS ▪ PRINTERS

Q. ***Where can you get the "Ultimate" SPARC® Workstation?*** *(See Answer below or call (800) 333-4834)*

A. *TDI, The Systems & Solutions Source.*

TDI has the ***experience*** to deliver quality workstations and the ***flexibility*** you need to select the ideal peripherals, disk, tape, memory, monitors and more. In addition to all this, **TDI** works with the industry leaders...***Axil Workstations, Micropolis, and Wangdat***. We will provide you with our ultimate service program ***"Total Customer Support"*** and guarantee ***Full Compatibility*** and ***100% SPARC® Compliancy***. We will provide custom configurations with a 48 hour turnaround time guaranteed, availability is stock. ***We make it easy!*** VAR's, Integrators, and OEM's give us a call today toll free at **1-800-333-4834.**

SATISFACTION GUARANTEED

We're the strength you need to stay competitive.

Axil WORKSTATIONS MICROPOLIS WangDAT™

TDI TECHNOLOGY DISTRIBUTORS, INC.
The Systems & Solutions Source

Circle 23 on Reader Inquiry Card

605 Neponset Street, Canton, MA 02021-1913
(800) 333-4834 • (617)821-2050 • Fax (617) 821-2185

SPARC® is a registered trademark of SPARC International, Inc.

5
Gigabytes

QIC
Data
Cartridge

GIGATEK
MEMORY SYSTEMS

1989 PALOMAR OAKS WAY, LA COSTA, CA 92009
TELEPHONE (619) 438-9010 FAX (619) 438-9120

Gigatek Introduces the First 5 Gigabyte QIC Data Cartridge and 0.5 Gigabyte QIC Minicartridge.

Gigatek is now shipping evaluation units of our 5 Gigabyte QIC data cartridge. In 1990 Gigatek introduced the first 525 megabyte QIC data cartridge. We thought 3 years was long enough to wait for a cartridge with 10 times the capacity.

But if 500 megabytes is all you need, now you will have two sizes to choose from - our 525 megabyte QIC data cartridge or our 500 megabyte (we like to think of it as 0.5 Gigabytes) QIC minicartridge.

Find out why companies that are household names in the quarter-inch cartridge business rely on Gigatek Memory Systems for the industry's most reliable QIC tape cartridges. Call our Customer Service department at (800) 669-1989 to put in your order for evaluation units of our 5 Gigabyte QIC data cartridge or our 0.5 Gigabyte QIC minicartridge.

Andrew K. Wrobel
Chairman, President
& Chief Executive Officer

0.5
Gigabytes

GIGATEK
MEMORY SYSTEMS

1-800
669-1989

Circle 24 on Reader Inquiry Card

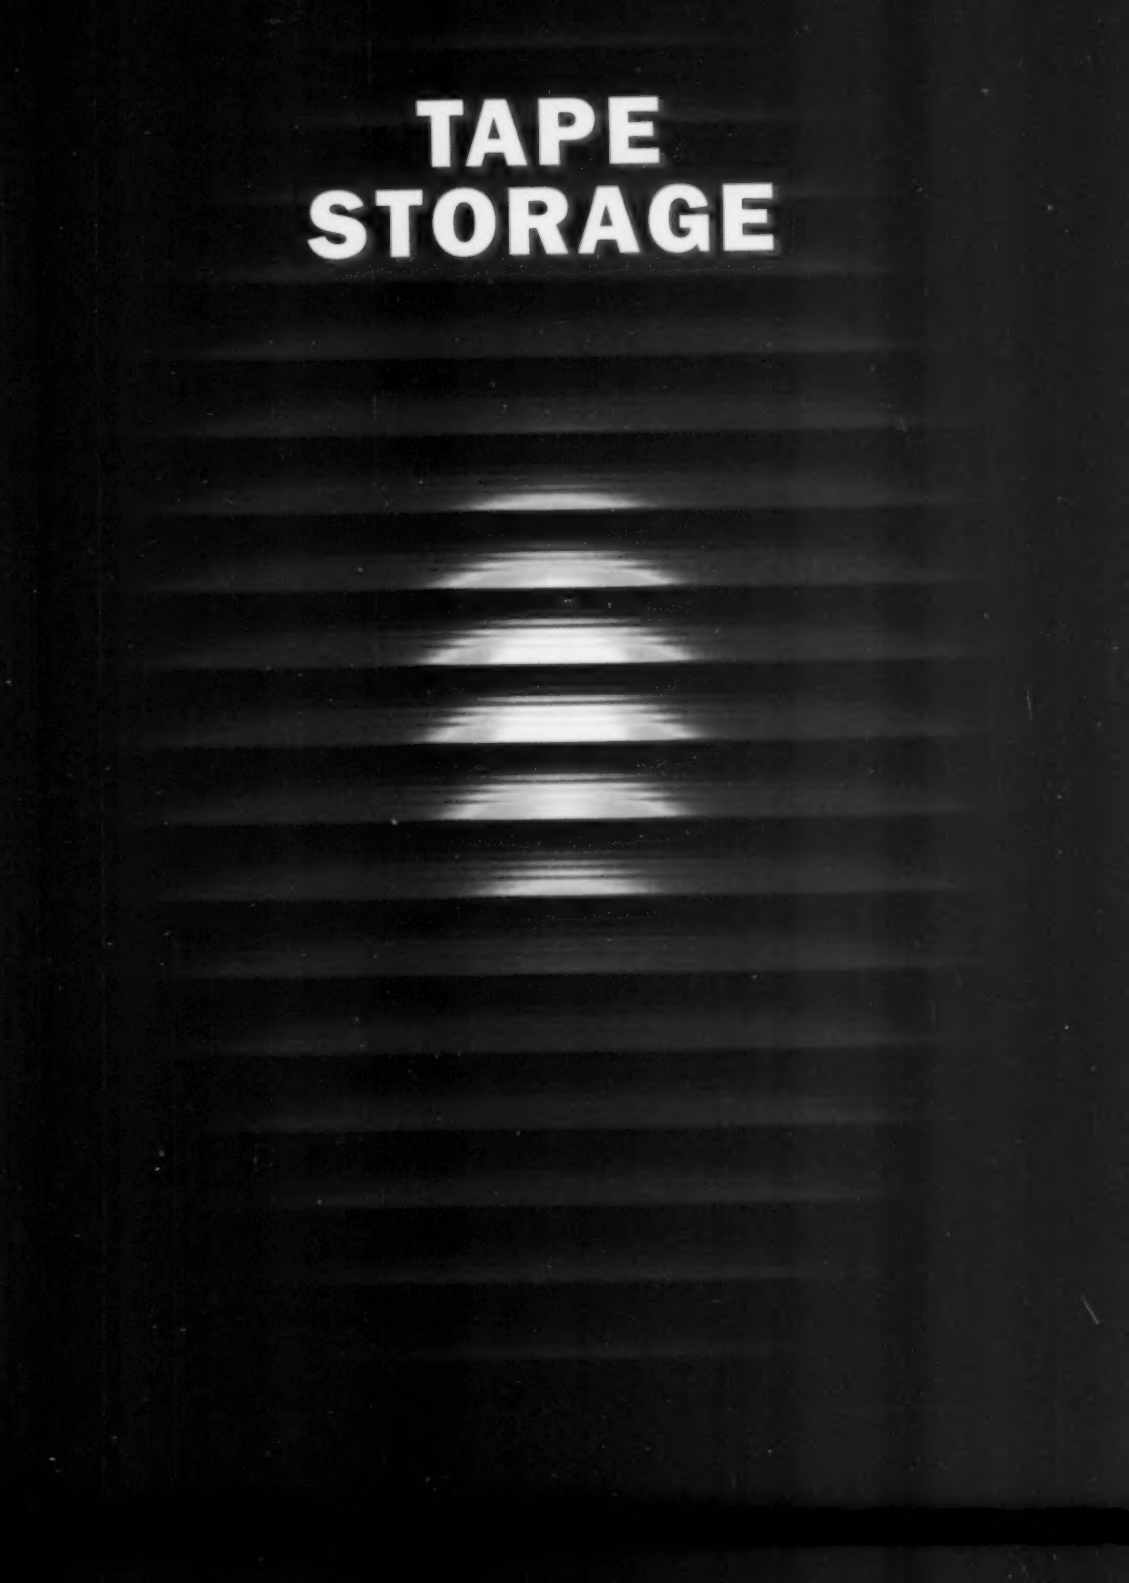
TAPE
STORAGE

Multi-User Networks Expand to Robotic Storage Systems

In all of the shifts and trends in data management today there is one constant: the increasing need for larger scale mass storage. It's an unfortunate fact that typically, soon after installation, the storage needs of multiple client/servers, LANS, WANS and network technologies quickly surpass their existing storage systems.

As these new networked systems continue to expand and multiply, they require more sophisticated storage systems, further expanded capacities, higher transfer rates, faster access times, smaller footprints and lights-out operation capabilities. And above all else, they require maximum cost-effectiveness.

IS Managers Turn To Robotics

No matter which combination of pieces of the puzzle are configured... servers, LANS, WANS, engineering workstations, desktop publishing systems, multi-user platforms... effective data management must include effective data storage. IS managers are looking past traditional individual workstation storage capabilities to multi-level robotic library storage systems to meet their needs.

Robotic storage systems are already delivering all of those solutions for large mid-range and mainframe systems. The downward translation to networked workstation systems can be made easily and efficiently. And the upgradability of robotic storage systems makes them the perfect storage partner for the new networks.

Incremental Expansion Flexibility

Upgradability affords unlimited storage system growth to match network system needs. Smaller systems can add modular desk-sized or rackmount automated systems that hold 100-500 Gigabytes and gradually grow into larger carousel systems.

Metrum Peripheral Products offers two

What to Drive.

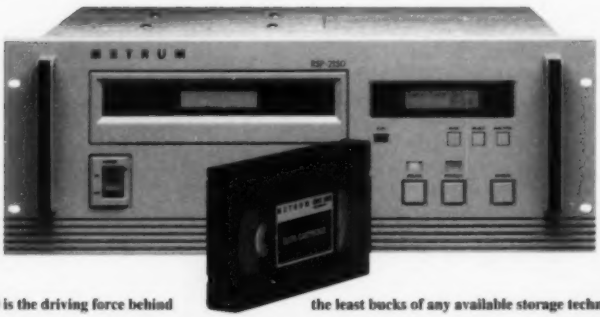

Metrum's RSP-2150 is the driving force behind the most cost-effective data storage. Using high density, helical scan technology, the drive stores up to 18 GB of information on a single half-inch METRUM ST-160 data cartridge. With a media storage cost of about 1/10¢ per MB it provides the most bytes for the least bucks of any available storage technology. But that's not all. The RSP-2150 backs up a 1 GB disk in 8.3 minutes with a data transfer burst rate of 4 MB per second, and a sustained rate of 2 MB per second. Which simply means that you can drive yourself home a little earlier each night.

Call 1-800-METRUM-2

For Product Literature

For technical information, call 303-773-4830
For reseller information, call 303-773-4574

For END-USER information circle 14
For OEM, VAR or SYSTEMS INTEGRATOR information circle 15

METRUM

NOBODY STORES IT BETTER.

automated data storage libraries for their half-inch tape cartridge drives. The RSS-48b holds 48 cartridges and stores from 696 to 864 Gbytes depending upon the capacity of the individual tape cartridges used. As data systems grow, it is possible to add additional RSS-48b units until it is more practical to step up to Metrum's RSS-600b, delivering 8.7 to 10.8 Terabytes on 600 tape cartridges.

Combined with state-of-the-art robotics, high-density digital tape drives assure high transfer rates and fast access to data files stored in the robotic libraries. Five Metrum RSP-2150 tape drives can be housed within the robotic systems. Each drive provides two Megabyte/sec. sustained transfer rate and a four Mbyte/sec. burst transfer rate through an industry-standard SCSI interface.

A plethora of new data management software products are beginning to take advantage of the high performance Metrum tape drives and robotic libraries to improve data availability, integrity and security. Continually upgraded drive technology and data management software assures that the storage systems will achieve greater and greater performance as the needs expand.

Metrum offers two capacity levels of half-inch tape cartridges, the ST-120 offering 14.5 Gbytes and the ST-160 with 18 Gbytes. The next generation cartridge, the ST-180, will store 20.5 Gbytes.

The new network systems themselves are adding to the storage systems' capabilities. Fiber optic technology is allowing for significantly enhanced throughput. Through fiber optic communications links, the total network can achieve increased speed, expanded data volume, improved quality and greater security. This combination adds even greater benefit to the existing advantages of robotic tape storage libraries.

The network systems software is also improving the overall storage capabilities, allowing for enhanced hierarchical data storage management. The final result is that the storage capacity is greatly expanded, and every operation in the entire network runs faster while achieving higher data integrity■

For END-USER information circle 10
For OEM, VAR or SYSTEMS INTEGRATOR information circle 11

Where to Park.

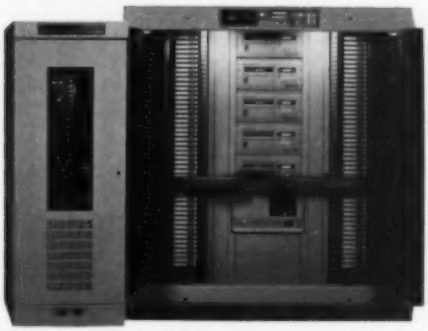

RSS-48b RSS-600b

If 18 GB isn't enough, just call for more back-up. The RSS-b storage libraries from METRUM offer the highest storage capacity per square foot than any competitor. Choose from two models. The RSS-48b which contains two tape drives and 48 cartridges for 864 GB of storage capacity. Or, the RSS-600b, which houses up to five tape drives and 600 cartridges for an incredible 10.8 TB of storage capacity. Each storage unit offers fast retrieval times, broadcast-proven precision robotic mechanisms, and inventory and security features. Either way, it's the perfect way to shift from drive to park.

Call 1-800-METRUM-2
For Product Literature

For technical information, call 303-773-4830
For reseller information, call 303-773-4574

For END-USER information circle 12
For OEM, VAR or SYSTEMS INTEGRATOR information circle 13

M ≡ T R U M

NOBODY STORES IT BETTER.

Virtual Tape Brings Cooperative Approach To LAN Backup

Virtual tape functions as a CPI-C/LU6.2 gateway between LAN backup software and the mainframe.

The push is on right now to find ways to link distributed and centralized computing environments in ways that leverage the strengths of both.

When it comes to LAN data backup, the mandates dictated by enterprise computing are clear and sensible. Maximize the usefulness of the corporate data center. Standardize the way data is protected for the entire corporation. Take advantage of IS expertise in backup and reliability of the mainframe. Utilize equipment that's already installed. All data is corporate data and should be managed and protected in a standard, reliable manner.

But acknowledging the advantages of backing up LAN data to the mainframe doesn't mean LAN administrators are ready to surrender control of their data to IS. Many are using server-based backup products from companies like Cheyenne, Legato, Palindrome and Sytos. While LAN administrators would welcome the prospect of reducing the headaches of tape rotation, off-site storage and data restoration, and the persisting doubts over reliability and disaster recovery, they are unwilling to give up the ease-of-use and control these server-based systems provide.

They don't have to.

Virtual Tape On The LAN.

Companies can choose to implement a cooperative approach to LAN data backup that appreciates both the way LANs work and the changing role of IS. Virtual tape on the LAN is a new solution that functions as a CPI-C/LU6.2 gateway between major LAN server backup products and dominant IBM MVS mainframe storage management utilities. It offers centralized storage and management of distributed data, while allowing the LAN backup process to continue to be administered and controlled from the LAN.

The virtual tape solution is open, interfacing with today's most widely installed LAN-based backup products so that current backup procedures are not disrupted. Instead of selecting a tape drive, administrators simply send data to the mainframe.

Intelligent Sharing

There's a clear need in the marketplace for an open solution like this. LAN managers want to retain administration over their data, and they're comfortable using their proven server-based backup systems. But they're tired of coping with the constant capacity management and media handling issues. IS managers are searching for ways to utilize the existing corporate data center resources without forcing mainframe-controlled products on the LANs. They're both under pressure from corporate management to maximize installed equipment investment. Virtual tape on the LAN fulfills all these requirements without disrupting their current procedures.

Glen McDermed, strategic planning analyst at Gartner Group says, "This kind of product will enable LAN-based backup products to use MVS as a virtual tape drive. This bridges the gap between local control and efficiency of centralized operations. It's the kind of solution that appeals to many organizations both strategically and politically. We see data centers becoming more

by Greg Coticchia
Emprise Technologies

WHY YOU NEED MAINFRAME CLASS TAPE DRIVES

EVEN IF YOU DON'T WORK WITH MAINFRAMES

NOT JUST FOR MAINFRAMES

The tape drive format you'll be building into your corporate computing systems for years to come is here: Mainframe Class Tape (MCT), the IBM 3480-compatible series of half-inch cartridge tape drives. It's the preferred medium for storage and interchange on mainframes and midrange computers worldwide.

What does MCT offer you over other tape drive technologies?

- Data transfer as fast as ten GB per hour native
- Trouble-free performance measured in years
- Guaranteed data interchange with mainframes worldwide

MCT™
MAINFRAME CLASS TAPE

The features that make MCT the tape drive technology of choice for the mainframes and midrange computers of today also make MCT ideal for the corporate computing systems of tomorrow.

MAKE WAY FOR THE NEW CORPORATE COMPUTER

Corporate computing systems are migrating to a new network architecture which provides distributed applications and databases to its users. This new network will include dedicated, powerful storage server workstations that will reliably manage corporate data assets. These workstations will demand the same performance and reliability from storage devices as are required by today's mainframe environment.

Large companies will migrate down from the traditional mainframe to the new network. Smaller companies will move up to the new network.

Even with all this change going on, one thing will remain constant: MCT.

WHY YOU NEED MCT

Corporate computing will still require mainframe class tape drives that transfer data quickly and can operate at up to 70% duty cycles for years without fail. Only MCT can provide these benefits.

Something else only MCT can provide: a format for exchanging data that everyone can use right now, regardless of platform. Even as MCT evolves from 18-track (3480) to 18-track/compression (3490) to 36-track/compression (3490E), MCT tape drives will be completely backward-read compatible, so your customers can continue using their existing tape libraries.

No other technology can guarantee that kind of interchangeability. Or data transfer rate. Or bullet-proof reliability.

That's why no other technology can be Mainframe Class Tape. Half-inch, IBM 3480-compatible cartridge tape drives. As you plan the corporate computing systems that will take your customers into the next century, make sure MCT half-inch tape drives are included.

They're your best guarantee that the systems you architect today will be up-to-date well into the future. With or without mainframes.

GET THE FULL STORY. CALL OVERLAND DATA TODAY AND SAY YOU'RE INTERESTED IN MCT.

WORLDWIDE *TEL: 619-571-5555* *FAX: 619-571-0982*
EUROPE *TEL: (+44) 734-891891* *FAX: (+44) 734-891897*

Overland Data Inc.

Circle 26 on Reader Inquiry Card

actively involved in LAN backup administration, and with good reason. LAN administrators are anxious to find a way to eliminate local tape handling and local tape drive purchases. They want to take advantage of IS resources including more reliable media and 24-hour operations staff available for tape mounts, but they are hesitant about losing local control over the backup process." "Utilizing the mainframe as a LAN backup data repository," McDermed continues, "has the potential to reduce the total cost of enterprise computing."

Many systems integrators and OEMs are interested in including this cooperative technology as part of their bundled product offerings. Major companies have already expressed interest in interfacing virtual tape on the LAN with high speed controllers, LAN backup packages and storage peripherals to create comprehensive systems.

The virtual tape solution conforms to the industry standard software and communications protocols most organizations already have in place. Operating as a NetWare Loadable Module (NLM) on the backup server or OS/2 process, it transmits data on the SAA for NetWare integration platform via CPI-C-based LU6.2 (APPC) to IBM MVS mainframes.

"Our customers have been demanding solutions that can leverage their existing IS staff and mainframe-based storage management tools to back up and archive their NetWare-based mission critical data," said Gerry Machi, vice president and general manager, NetWare Systems GroupEnterprise Products, Novell, Inc. "The virtual tape solution, along with NetWare for SAA, provides NetWare customers with a comprehensive, enterprise-wide solution for their storage management needs."

Fitting into the Distributed Environment

Virtual tape provides a gateway between LAN backup products on the NetWare server and centralized storage resources on the mainframe. It comes in two parts: one runs as a NetWare Loadable Module (NLM) on a NetWare LAN server; the other operates as a VTAM application under MVS on an IBM mainframe.

To users of the LAN backup products, the mainframe is transparent. A LAN administrator can send backup data to the data center just as if it were another local tape drive. The mainframe seems like a tape with an almost unlimited capacity. Optionally, the LAN backup data, as it ages, can move automatically through a hierarchy of mainframe direct access storage, tape and even off-site storage, if the mainframe environment includes a data migration facility, such as IBM's DFHSM.

Running on the same server, the NLM receives data from the LAN backup product just like a SCSI-attached tape changer (Fig 1). The virtual tape transmits the LAN backup data to the mainframe across an LU6.2 connection defined through an SAA and X/Open standard CPI-C. The data from each LAN backup session resides in a VSAM session file on a mainframe direct access storage device (DASD). Each VSAM file appears to the LAN backup product like a backup session on tape. Each of these "tape sessions" are grouped into logical tape volumes and tracked through a control file on the mainframe. If data management and migration facilities like DFSMS and DFHSM are in use, the VSAM session files can be migrated automatically through the hierarchy of mainframe storage resources to slower DASD, to local tape and to archive tape for off-site storage.

Unlike many distributed storage management products, the virtual tape solution capitalizes on existing LAN backup solutions, while providing maximum flexibility of configuration between LAN and mainframe resources.

By interfacing with server-based backup products, this technology preserves the benefits they provide and extends their capabilities even further. It doesn't disrupt backup procedures. It offers another choice of device. Instead of selecting an automatic tape changer, administrators send data to the mainframe.

The virtual tape solution reduces costs and increases interoperability between LAN and mainframe domains in the distributed environment by:

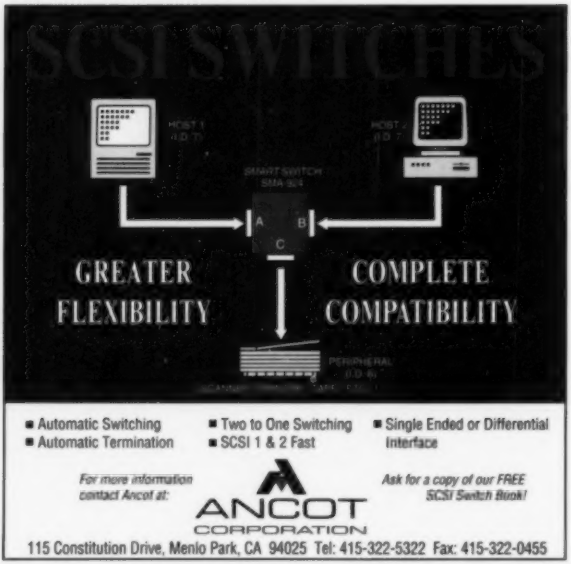
SCSI SWITCHES
GREATER FLEXIBILITY
COMPLETE COMPATIBILITY
A
B
C
■ Automatic Switching
■ Automatic Termination
■ Two to One Switching
■ SCSI 1 & 2 Fast
■ Single Ended or Differential Interface
For more information contact Ancot at:
ANCOT CORPORATION
Ask for a copy of our FREE SCSI Switch Book!
115 Constitution Drive, Menlo Park, CA 94025 Tel: 415-322-5322 Fax: 415-322-0455

Circle 27 on Reader Inquiry Card

Around the world over nine million businesses share a common practice. They back up and store important data on quarter-inch data cartridges. And the brand they choose most is 3M. For good reason.

3M data cartridges are built for reliability—from the tape formulation to the cartridge itself. Whatever capacity you need—formatted or unformatted—3M cartridges are ready to save the day with proven performance. That's why more businesses protect important information on 3M brand data storage products than any other brand in the world. For more information and your free "Data Security Handbook" call 1-800-888-1889, ext. 43.

The common denominator of business around the world.

3M data cartridges require compatible drives.

Circle 28 on Reader Inquiry Card.

3M *Reliability*

LAN Backup

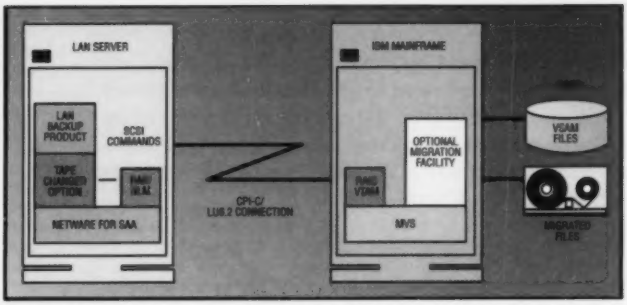

Fig 1 The LAN administrator sends data to the NLM, just as if it was another tape drive.

- preserving LAN and IS installed hardware and software investment;
- minimizing manual media handling;
- eliminating issues of media reliability and capacity management;
- maintaining LAN administrator control;
- providing IS with centralized management of corporate data;
- optimizing existing mainframe storage processes;
- extending the usefulness of trained IS staff; and
- capitalizing on existing, proven products specifically designed for LAN backup.

LAN backup administrators can utilize the virtual tape on the mainframe to serve whatever backup, migration or disaster recovery needs they choose. Selection of files, file groups or directories and scheduling of backup, restore or migration processes is accomplished using the normal, familiar and easy-to-use features of proven, industry standard LAN backup products. Any resource normally backed up by the server-based LAN backup product may be directed to the mainframe.

At the same time, administrators retain the flexibility to continue using their LAN backup products of choice. They incur no new training costs and sacrifice none of their ability to direct the backup process.

Centralized Vs. Distributed Control

No two organizations have a common approach to distributed computing. Virtual tape on the LAN adapts to meet the requirements of each customer's specific installation, whether it be strictly centralized, predominantly distributed, or something in between.

Virtual tape can be implemented to the degree and in the manner required by individual installations and still give organizations the freedom to alter their LAN backup procedure as their environments change (Fig 2).

The following are just three examples of how virtual tape can be used to enhance backup and disaster recovery functions in specific, distinct environments. Many more options are possible.

- **Example #1: Distributed Administration** Use virtual tape only as a long term archival solution. LAN administrators continue to perform normal data backup to local media. Mission critical files are automatically sent to the mainframe for archiving.
- **Example #2: Centralized Administration** Use virtual tape only as a centralized solution. All LAN servers are unattended. Complete backup administration is executed from the data center's standard LAN backup console. Local media handling is completely eliminated. Remote servers run backups unattended. All corporate data is secured in a consistent manner in a central location.
- **Example #3: Cooperative Administration** Use virtual tape as a cooperative solution for LAN backup to the mainframe. LAN administrators continue to perform short-term local backups. Files older than a specified number of days are automatically sent to the mainframe for safe, centralized storage and disaster recovery. IS provides proven, centralized storage management.

Optimizing LAN Backup

Virtual tape on the LAN offers unique features that optimize LAN backup operations:

Open system architecture based on industry standards. Functioning as an NLM on the backup server, data is transmitted via CPI-C-based LU6.2 to IBM MVS mainframes.

Seamless interface with installed server-based LAN backup products. Virtual tape provides the appearance of a server attached tape changer device.

LZW data compression. Employing LZW compressor technology, virtual tape automatically compresses data as it is shipped to the mainframe for optimized transmission rate.

VSAM file maintenance. Each "tape" session is maintained as an individual VSAM file. Sessions are grouped into logical tape volumes referenced by a control file on the mainframe.

Archiving. VSAM session files are eligible to be migrated to mainframe hierarchical storage under the control of IBM's DFSMS, DFHSM, or other popular independent software vendor file migration facilities.

Volume management. The creation and deletion of logical tape volumes is at the discretion of the LAN backup product and administrator.

LAN-based or mainframe activity log. All activity is logged (to SYSOUT or a file) on both the LAN and the mainframe. The log may be displayed and printed.

Mainframe reporting. A facility on the mainframe can be used to report on the contents and status of the mainframe backup repository, including which files from various LAN systems were backed up and when, and what mainframe resources are in use providing backup of LAN data.

Virus protection. Virtual tape can include popular virus protection software for system-wide security.

Unattended operations. Through the use of central command facilities, depending on the LAN vendor, all administration may be centralized. For sites where local administration is not possible, each product can be run automatically with no local interaction for backup.

Error recovery and handling. In the

Instant Rebates

TECMAR

Get an instant rebate of $25 to $500 when you buy selected Rexon family backup tape drives.

Tecmar Backup Solutions

Bundled products for NetWare, DOS and OS/2		Rebate
QICVault 720is	720MB internal, FlexPAK* software	$50
QICVault 720ex	720MB external, FlexPAK* software	$50
DATaVault 2000	2GB DAT external, FlexPAK* software	$50
DATaVault 4000	4GB DAT external, FlexPAK* software	$100
ProLine DAT 2000	2GB DAT external, ProServe 4.0 NLM software	$150
ProLine DAT 4000	4GB DAT external, ProServe 4.0 NLM software	$200
ProStack 48	48GB DAT autoloader, ProServe 4.0 NLM software	$500
Macintosh® Solutions		
QT 250 MAC Bundle		$25
QT 525 MAC Bundle		$25
DATaVault 2000 MAC Bundle		$25
DATaVault 4000 MAC Bundle		$25

*FlexPAK includes backup software for DOS® (QTOS), OS/2® (Sytos Plus for OS/2) and Novell® NetWare® (QTOS NLM).

Wangtek "Bare" QIC Drives

For UNIX and integration		Rebate
5150ES	150MB internal	$25
5525ES	525MB internal	$25
E5525ES	525MB external	$50
51000HT	1.2GB external high transfer	$50

WANGTEK

WangDAT

WangDAT "Bare" DAT Drives

For UNIX and integration		Rebate
3100	2-4GB internal	$50
3200	4-8GB internal	$75
3200SE	4-8GB external	$100

Instant Rebates are available February 1 through March 31.
To order, contact your distributor. Or call the backup information line:

1-800-4-BACKUP

REXON
SOFTWARE ■ STORAGE ■ SYSTEMS
Sytron ■ Wangtek ■ WangDAT ■ Tecmar

INGRAM MICRO
1 (800) 456-8000

MERISEL
DISTRIBUTED BY
1-800-MERISEL

Tech Data CORPORATION
1-800-237-8931

GATES/FA DISTRIBUTING
800-332-2222

All Tecmar product names are trademarks or registered trademarks of Rexon/Tecmar, a subsidiary of Rexon, Inc. Other product and brand names are trademarks or registered trademarks of their respective companies.

Circle 29 on Reader Inquiry Card

LAN Backup

event of communication failure during backup, the mainframe VSAM session file is closed and the LAN backup product is notified immediately. Depending on the functionality of the specific LAN product in use, the backup process either begins again or is resumed just prior to the point of failure without data loss.

Fast local restore with dual write. Users have the option of backing up data both to a locally attached tape drive and to the mainframe, in succession. This provides the maximum level of backup flexibility, letting users choose between the speed of restoring files from tape or, when tapes are stored off-site, the convenience of data recovery from the mainframe.

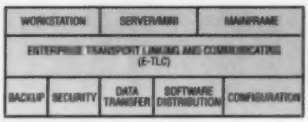

Fig 2 E-TLC provides the inter-system management link between backup software and the mainframe.

Windows- or DOS-based administration. Depending on which LAN backup product is installed, all administration can be achieved using the most popular and predominant operating systems, Windows or DOS.

A Tactical Solution To A Strategic Problem

"With the introduction of virtual tape on the LAN, IS is finally able to access LAN data without depriving LAN administrators of control over their own information," says Michael Peterson, industry analyst and president of Peripheral Strategies. "After all, IS and LAN managers both have the same overriding objective: to protect critical corporate data." ■

Greg Coticchia *is the general manager of Emprise Technologies (Bridgeville, PA). He completed his undergraduate degree in Industrial Engineering, followed by an MBA, at the University of Pittsburgh.*

For more information circle 109.

2 Year On-Site Warranty

By definition there can be only one "BEST"

... and in the world of data interchange the logo always says **"M4 Data"**

WITH OUR PEDIGREE OF MORE THAN TWENTY YEARS EXPERIENCE IN DESIGNING AND PRODUCING 9-TRACK TAPE DRIVES FOR HIGH-PERFORMANCE, HIGH-DUTY CYCLE APPLICATIONS, YOU CAN BE ASSURED OF UNCOMPROMISED QUALITY WITH M4 DATA. FROM OUR 150 IPS QUAD DENSITY 9914 PRODUCTS TO OUR NEW BRIEFCASE-SIZED 9906 WE HAVE THE BEST SOLUTION FOR YOUR DATA INTERCHANGE NEEDS. NO MATTER IF IT IS PERSONAL COMPUTERS, HIGH PERFORMANCE WORKSTATIONS, MID-RANGE OR MAINFRAME ATTACHMENTS, M4 DATA CAN MAKE IT HAPPEN. YOU CAN SETTLE FOR GOOD, EVEN FIND BETTER, BUT THERE IS NO OTHER BEST BUT M4 DATA.

...for more information, call us at

1-800-240-8270

M4 Data Inc.

M4 Data Inc.
3815 N. U.S. 1, Bldg. 118, Cocoa, Fl 32926
Phone: (407) 639-6487 Fax: (407) 639-9800

M4 Data Ltd.
10855 Sorrento Valley Road, Ste. 201A, San Diego, CA 92121
Phone: (619) 597-0516 Fax: (619) 597-0936

Circle 30 on Reader Inquiry Card

SONY

SONY PRESENTS AN OPEN AND SHUT CASE.

The evidence is clear. Sony Enhanced Data Grade D8 cartridges deliver unsurpassed performance. The proof? A redesigned lid that withstands over 10,000 openings and closings. An outer shell and media resistant to both heat and shock. Plus an archival life exceeding 30 years. And up to 1,500 full back-ups. Hear all the facts. Call 201-476-8199.

©Sony Electronics Inc. Reproduction in whole or in part without written permission is prohibited. All rights reserved. Sony is a trademark of Sony.

Circle 31 on Reader Inquiry Card

The industry is applauding the orchestration of a powerful new technology from the leader in tape storage, Exabyte.

And Vice Versa.

The Exabyte EXB-2502 Minicartridge Tape Subsystem is the ideal SCSI drive for PC, workstation, and LAN applications. This minicartridge tape drive (DC 2000) stores an unprecedented 2 gigabytes of native data in half the time of any other 3.5-inch tape drive. And reliability is evident through a 200,000 hour MTBF.

The 3M minicartridge is the first to use advanced barium ferrite technology. Its durability and environmental stability make it an ideal storage medium for backup and archiving. And it provides the best migration path for higher quarter-inch data densities.

Together, they will bring harmony to all of your data storage needs. Call your Exabyte representative today or 1-800-EXABYTE for special OEM information.

The Exabyte EXB-2502 Minicartridge Tape Subsystem and the barium ferrite minicartridge from 3M.

EXABYTE®

We're Backing It Up.

©1994 Exabyte Corporation
Exabyte is a registered trademark of Exabyte Corporation.
3M is a registered trademark of 3M.

Circle 32 on Reader Inquiry Card

SANKYO

Sankyo has established a reputation for producing the industry's highest quality 1/4-in data cartridge tape drives by leveraging *a vertical integration policy* with traditional *quality by design.*

Sankyo Tape Products Division develops, manufactures, and currently offers a full line of industry standard DC-6000-class data cartridge tape systems in 5-1/4-inch, half-height form factors. Products are available in a variety of configurations as either board drives for OEMs or complete subsystem solutions. Every unit offers field-proven reliability backed by a **two-year tape drive warranty.**

Sankyo fully supports QIC standards to enable data interchange and multiple sourcing. Sankyo continues its commitment and investment in future QIC development standards for next generation technologies.

SANKYO. . .CERTIFIED

Sankyo products are software certified for Cheyenne *ARCserve;* Sytron *Sytos Plus;* NovaStor *NovaBack;* Microlite CTAR UNIX; Emeritus *TapeWare;* and Dantz *Retrospect* and Surf City *SurfGuard.* Sankyo products supported by Adaptec, Always Technology, Buslogic, DPT, Future Domain and UltraStore SCSI host adapters.

SANKYO . . . For VARs and SYSTEM INTEGRATORS

For VARs and System Integrators,Sankyo means strategic partnerships enabling you to obtain compatible QIC technologies, *Quality by Design, Just-in-Time (JIT) Deliveries* and Applications Support from a service oriented organization.

Sankyo . . . Quality By Design

Tape Products Division

Corporate Offices:
2649 Campus Drive
Irvine, CA 92715
Phone: (714) 724-1505
Fax: (714) 724-1501

Europe:
Unit 9 Eghams Court
Boston Drive, Bourne End
Bucks SL8 5XS
United Kingdom
Phone: (44) 628-810260
Fax: (44) 628-819435

Eastern Region:
498 Palm Springs Drive, Suite 100
Altamonte Springs, FL 32701
Phone: (407) 339-7940
Fax: (407) 260-2121

Call: 1 (800) 4-SANKYO
[1 (800) 472-6596]

Distributed by:
1 (800) 227-1693

Pioneer

Circle 33 on Reader Inquiry Card

Defining Fibre Channel For Network And Channel Protocols

Fibre Channel is a new I/O interface for network and channel users to expand interconnects for tape, disk and other peripherals.

Fibre Channel is the general name of several standards being developed by the American National Standards Institute (ANSI) which together define new protocols for flexible information transfer. It began in 1988 as an extension of work on the Intelligent Peripheral Interface (IPI) Enhanced Physical standard and branched out in several directions since then. The primary objectives are:

- to support multiple physical interface types;
- to provide a means to interconnect those various interface types;
- to provide high-speed transfer of large amounts of information; and
- to provide relief for the growing set of physical interfaces from which manufacturers must choose, many of which have limited general use.

Channels vs. Networks

There are two basic peripheral protocols for device communication: channels and networks. Traditionally, the term channel refers to a peripheral I/O interface (to a host computer) that transports large amounts of data between devices. The processing overhead is kept to a minimum by handling data transfer in hardware with little or no software involvement once an I/O operation begins. On the other hand, the term network refers to an I/O interface that usually supports many small transactions with higher overheads, usually the result of software involvement in the flow of information.

Channels usually operate in a closed, structured, predictable environment where all devices that can communicate with a host are known in advance and any change requires host software or configuration table changes. This advance knowledge of the configuration is very important in the high performance levels achieved by most channels.

Networks, on the other hand, operate in an open, unstructured and essentially unpredictable environment. Almost any host or device can communicate with any other device at any time. This unpredictable nature requires more software support for verification, setting up sessions, and routing transactions to the correct software service.

Fibre Channel attempts to combine the best of these two opposing methods of communication into a new I/O interface that meets the needs of both channel users and network users. It virtually sidesteps the issue by providing a means to transfer data between one buffer at the source device and another buffer at the destination device. All channel and network protocols require buffers to hold data they send and receive.

Fibre Channel is unaware of the content or meaning of the information being transferred. It increases connectivity from tens of devices to thousands or even hundreds of thousands of devices, increases the distance between devices, and increases the transfer rate by four to five times over the most popular channels and up to 100 times over popular networks.

Fibre Channel does not incorporate a command set as SCSI or IPI do, but it does provide a mechanism to define protocols that use the common information transfer functions of Fibre Channel to accomplish either

by Jan Dedek and Gary Stevens
Ancot Corp.

of those command sets. This means that significant software investments in various command sets from older I/O interfaces can be directly transferred to Fibre Channel. This separation of I/O operations from the physical I/O interface is an important feature of Fibre Channel and makes it possible to use multiple command sets simultaneously.

The Fabric Of Fibre Channel

Fibre Channel devices are called nodes, each of which has at least one port to provide access to the outside world (i.e., other nodes). The components that connect two or more nodes together are called a topology. All Fibre Channel systems will have only these two elements: nodes and topologies. The Fibre Channel standard defines three topologies, but the main emphasis has been on one called the Fabric.

Each Fibre Channel port connects to a pair of fibres—one to carry information into the port and one to carry information out of the port. In Fibre Channel, fibres are wires or optical strands. This pair of fibres is called a link and is part of each topology (Fig 1). Data is always transmitted in units called frames over these links.

A Fabric (Fig 3) permits dynamic interconnections between nodes through ports connected to the Fabric. Each port in a node, called an N-Port, is attached to the Fabric through a link. It is possible for any node to communicate with any other node connected to the same Fabric using the services of the Fabric.

The function of the Fabric is similar to that of a telephone system—we dial some number, the telephone system finds a path to our destination, the phone rings, and the other party answers. If a switch or link goes down, the telephone company reroutes calls on other paths and the callers seldom notice. Most of us are not aware of the intermediate connections that the telephone company switches make to complete our "simple" telephone call.

A Fabric is of similar complexity to the telephone switching devices. In Figure 3, if Node A needs to talk to Node C, the information is first sent to the Fabric at Port a. The Fabric makes an internal connection or set of connections to Port c. The information is then sent to Node C. It may be necessary to select several paths internal to the Fabric before arriving at Port c. A Fabric is often called a switched

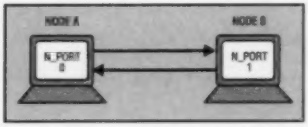

Fig 1 A point-to-point topology links only two ports.

topology or a cross-point switch topology.

When the topology is separate from the nodes, new technologies with new speeds and new functions may be introduced without making all previous investment in your nodes obsolete. The Fibre Channel can accommodate a mix of connections with different speeds or functions.

Besides the Fabric topology, Fibre Channel standard defines two other topologies: one is called Point-to-point (Fig 1) with only two ports connected together. The other topology is called Arbitrated Loop (Fig 2), which is a relatively

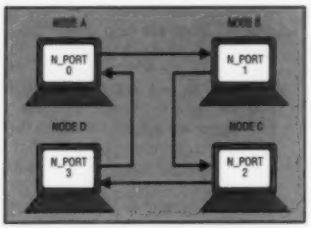

Fig 2 The Arbitrated Loop is a cost-conscious alternative for connecting numerous ports.

simple and low cost version for connecting smaller number of ports (a few tens) together in a loop fashion.

To better understand the function of Fabric topology, let's continue the analogy with the telephone system—we dial a destination number. We don't have to know the precise route to the phone when we call. The routing is taken care of by the telephone system. The Fabric of Fibre Channel performs the same function—you give a destination ID in the header preceding the actual data and the Fabric routes frames to that destination. If you dial a number incorrectly, the telephone company tells you that the number is not in service. The Fabric rejects frames for invalid destinations in much the same manner. If there is a congestion problem en route, the fabric responds with a busy signal and the calling port tries again.

Just as the telephone company can and does configure many routes between various points to provide reliable service, a Fabric can have multiple paths between switches to carry more traffic and to provide backup paths in case of a switch or link failure.

How Many Ports Can There Be?

The Fabric is constrained only by the number of node ports, N-Ports that can be identified in the destination address field of frames. That limit is set at a little over 16 million. For a single integrated system, this should meet all needs for quite a while.

The Fabric and Arbitrated Loop topologies of Fibre Channel can be intermixed in one system to provide a variety of service and performance levels to the nodes. In addition, a Fabric may use other networks such as SONET or ATM over SONET to achieve longer distances than can be achieved by the links to N-Ports. These links would be found be-tween switches distributed over a large geographic area.

Functional Levels

Some topics just naturally belong together and so it is with Fibre Channel. Items dealing with producing reliable and testable optical links are quite removed from how to recover from a missing frame of a sequence. These separate areas of interest are called functional levels in the Fibre Channel standard. There are five levels defined and each one is labeled FC-X.

■FC-0 defines the physical portions of Fibre Channel, including the media types, connectors, electrical and optical characteristics needed to connect ports together. This level is in the FC-PH standard.

DAT

High Capacity 2, 4, and 8GB DAT Solutions!

Compatibility

FREE Software Offer*... Up to $295 value!

Buy a 2GB WangDAT model 3100 or 4GB model 3200 drive in November/December '93 and choose one of the following items, absolutely ***Free!******

- ❑ *Sytos Plus® for DOS*
- ❑ *Sytos Plus for OS/2®*
- ❑ *Server Backup plus™ NLM*
- ❑ *Two 90m DAT Cartridges*

Reliability

Industry Leadership

Available Now!

High Capacity, Performance DAT Storage Solutions
Our full family of WangDAT drives includes 2, 4, and 8GB models, highlighted by our top of the line 3400DX, an 8GB DDS-2 data compression drive. WangDAT's drives combine unmatched compatibility, reliability, and mechanical design.

Cross Platform Compatibility for Increased Functionality provides plug and play operation on Novell networks, UNIX systems as well as easy integration with the leading tape software for Apple Macintosh, DOS and OS/2 based PCs.

Improved Reliability through Superior Design Features
WangDAT offers drives with an MTBF up to 180,000 hours based on a rugged 30% duty cycle. Our *DATAClass®* mechanism provides gentle tape handling for increased tape longevity, and with three levels of ECC we deliver unparalleled data integrity.

WangDAT is the Unsurpassed Leader in DAT Technology.
We pioneered DAT in the industry with the first half-height DAT drive. Today, WangDAT delivers the broadest line of DAT products to meet your needs.

Call your WangDAT distributor today or call us at **1-800-992-9916** for complete details on this great offer!

ISO 9001 REGISTERED FIRM

WangDAT
PRODUCTS BY REXON

*Free offer limited to one of the listed items per WangDAT Model 3100 or 3200 drive purchased from an *Authorized WangDAT U.S. Distributor* between Nov. 1st, 1993 and Dec. 31st, 1993. Proof of purchase required. Offer expires Feb. 28th, 1994. ©1993 Rexon Inc. All rights reserved. WangDAT and the "Slant D" are registered trademarks of WangDAT Inc. All other marks are the properties of their respective holders.

Circle 34 on Reader Inquiry Card

■FC-1 defines the transmission protocol, including the 8B/10B encoding, order of word transmission, and error detection. This level is in the FC-PH standard.
■FC-2 defines the signaling and framing protocol, including frame layout, frame header content and rules for use. The bulk of the FC-PH standard is for this level.
■FC-3 defines common services that may be available across multiple ports in a node. This level has no standard at present.
■FC-4 defines the mapping between the lower levels of Fibre Channel, above, and the command sets that want to use Fibre Channel. Here you will find mappings of SCSI, IPI-3, HIPPI, etc. Each of these command sets or existing interfaces is given a separate standard since the interested parties don't need or want to be bothered with extraneous information about other mappings. If you are doing SCSI, you don't really care about IPI-3 or HIPPI.

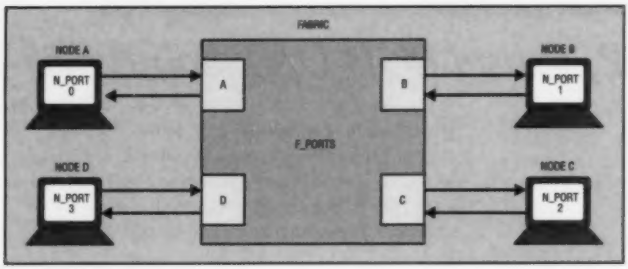

Fig 3 A fabric topology provides an architecture for interconnection in a more dynamic topology.

Classes Of Service

Class 1: Direct contact, a face-to-face meeting. In Fibre Channel, this is a hard or circuit-switched connection, functioning similar to a dedicated physical channel. Using this class, the path is not available to other hosts, and is nicknamed the "selfish" channel.

Class 2: Notes, messages, and short telephone calls. The Fibre Channel analog is a connectionless, frame-switched link. It carries less delay than class one, and no uncertainty of delivery. If there is congestion, a busy signal is returned. The path is not dedicated in order to more efficiently use the link's bandwidth.

Class 3: Shouting down the hall or an advertisement. Known as the "hopeful" link, this class is a connectionless service that allows data to be sent rapidly to multiple devices attached to the fabric, but no confirmation of receipt is given.

In Fibre Channel there is an optional mode called Intermix. Intermix reserves the full Fibre Channel bandwidth for a dedicated conection, but also allows connectionless traffic to share the link if there is bandwidth available.

Performance

A basic premise of Fibre Channel is that it is able to replace serial HIPPI. Although the parallel version does not allow very long distance, it is fast—100 Mbytes/sec. Also, IPI-3 at 25 Mbytes/sec is a second performance mark for Fibre Channel. These two speeds as well as 50 and 12.5 Mbytes/sec make up the four speed choices for Fibre Channel (100, 50, 25, and 12.5 Mbytes/sec). To support a net rate of 100 Mbytes/sec, we must have a bandwidth of ((800 ÷ 200) ÷ 62.5) or 1062.5 Mbaud to cover the encoding and framing protocol overhead.

Since the calculation is linear, we just divide by 2, 4 or 8 to arrive at the corresponding rates for the lower speeds. The complete set of rates for Fibre Channel is 1062.5, 531.25, 265.625, and 132.8125 Mbaud.

The low overhead of Fibre Channel is possible partly because of the hardware intensive requirements. Each port must continuously send frames or "Idles" on its outbound fibre. Each port in a system receives a continuous stream of bits continu-ously at over a Gbaud (1.0625 Gbaud to be exact); there is no time to make decisions in any type of processor. There must be a real-time hardware assist in each port if the kind of data rate advertised is to be achieved.

Overhead is a fixed cost or propor-tional cost that cannot be eliminated. In Fibre Channel, the overhead is measured in lost time (increased latency) or bytes per second of transfer rate lost. The 8B/10B encoding is the largest component of the overhead at 25%, but the increased baud rate compensates for it. Additionally, the increased baud rate compensates for framing protocol. The effective rate, in Mbytes per second, is about 75% of the real transfer rate in megabaud per second. That is, the electronics and optics speeds were increased to achieve the desired net transfer rate.

When HIPPI uses class 1 or class 3, it can get more than 103 Mbytes per second over a single fibre. If HIPPI uses class 2, the effective transfer rate is just over 100 Mbytes per second, but Fibre Channel meets the requirement we stated earlier in any class of service.

Pick Your Media

Now that we have the speeds established, and the rationale for those speeds, it is time to discuss the media choices available in Fibre Channel. Interfaces are driven optically and electrically. The two types can be combined in a single system when there is a Fabric or other media-type converter available; a Fabric is the only speed converter available in Fibre Channel. This is important to system integrators. There is no need to specify that a system with a Fabric has all the same type of media, lasers, distance or speed. A mix-and-match system is quite reasonable and usually less expensive that a homogeneous system, and the performance should meet all requirements.

Interoperability may seem like an issue here, but a Fabric is normally required for larger systems. If the

A Case For QIC

Fact 1

The Wangtek 9000 Series of Thin-film QIC Drives Adds New Meaning to Storage Systems Performance and Capacity:

Superior performance over 8mm solutions, *at half the price.*
2.1, 5, and 10GB capacities, thanks to our thin film/MR head technology.
Features: Auto *Throttle*™ for dynamically variable transfer rates from 300-778 KBS, and SOFT Load® automatic cartridge loading for VCR-like ease of use.

Fact 2

The 9000 Series Gives You 100% Backward Compatibility With Your Installed Base:

Read compatibility with all major QIC formats protects your data investment now and in the future.
SCSI I and SCSI II interface options provide host adapter compatibility.
Reliability and low cost per megabyte provide compatibility with your budget.

Fact 3

QIC Technology Ensures a Future of Savings:

We're committed to QIC market leadership.
Our leading edge technology will advance QIC to higher capacities and performance.
Wangtek 9000 Series QIC storage solutions are available today!

Solution

Call Wangtek *Now* at 1-800-992-9916 for Product and Sales Information.

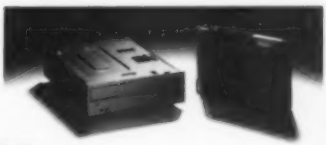

ISO 9001
REGISTERED FIRM

WANGTEK™
PRODUCTS BY REXON

©1993 Rexon Inc. All rights reserved. Wangtek is as trademark and the Circle "W" is a registered trademark of Wangtek Inc. All other marks are the properties of their respective holders.

Circle 35 on Reader Inquiry Card

Fabric vendors do as we hope, the Fabric ports can be changed easily from one type to another with little or no interruption to the rest of the Fabric.

The Future Of Fibre Channel

Before a standard has a future, it must have a present. Fibre Channel is here; it is receiving considerable attention from most high-tech corporations. Some components, expected in the first half of 1994, will make viable systems a reality. Since the time to produce a new product is considerable, especially with new technologies and interfaces, the plans for these products have been in place for some time.

All of the components to build a Fibre Channel system are available now. Some of these components were built based on an earlier level of the standard than is now available, but valuable development work can be achieved with these components. These systems can operate at either the 25 Mbytes/sec rate or the 100 Mbytes/sec rate. These are baseline feasibility systems, or beta test systems as they are sometimes called.

There are plenty of suppliers of Fibre Channel components. Almost any high-tech corporation has parts, or is buying parts or is planning to buy parts for Fibre Channel.

The first beta test site for Fibre Channel is at Lawrence Livermore National Laboratory in California. The laboratory has 32 by 32 crosspoint switches, nodes, ports, and software which are working on several UNIX platforms.

The ANSI committee X3T9.3 is responsible for developing the suite of Fibre Channel standards. Open meetings, with representatives from many companies, are held at regular intervals. This committee produces the standard, but it is not responsible for making judgments about or for compliance testing of implementations. Compliance or conformance testing is usually done by independent companies or organizations.

The Fibre Channel Association (FCA) is a cooperative organization independent of the standards committee. This is a fee-for-joining organization that is open to any company. Its purpose is to promote the use of Fibre Channel.

The Fibre Channel Systems Initiative (FCSI) is a closed organization made up of representatives from Hewlett-Packard Company, IBM Corporation, and Sun Microsystems Corporation. Its purpose is to promote early implementation of Fibre Channel for their workstation products. By defining interoperable subsets of the standard, products from the three companies can interoperate. Also, suppliers of Fibre Channel products to these companies will be required to meet the subset requirements. ■

Jan Dedek *is president of Ancot Corp. (Menlo Park, CA)* **Gary Stephens** *is president of FSI consulting (Tuscon , AZ).*

For more information circle 118.

UNLIKE THE OTHERS, OUR TAPE DRIVE WON'T DO THIS.

If you've ever had visions of your tape being eaten, torn or mutilated, here's something to digest.

Philips LMS drives use fixed heads with air bearings rather than the technology used in a regular VCR. Instead of the tape wrapping around a rotating head, it floats on air. There's absolutely no contact. And, therefore, no wear and tear.

We also offer 16.8 GB capacity. A quick data transfer rate of 6MB/second. Outstanding performance. And reliability you can't expect from any other drive.

For more information and a White Paper on why you need 1/2-inch tape technology, please call 1-800-777-5674.

PHILIPS
L M S

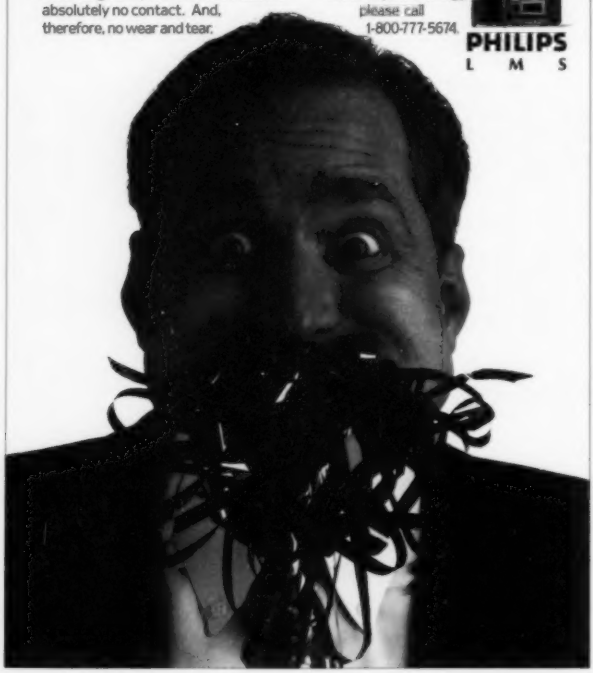

Circle 36: Call me, I'm interested.
Circle 37: Please send literature.

Library Integrators Address Key Technical Issues In Automating Network Storage

Dan Stromska
Advanced Digital Information Corporation

Robotic autochangers are one solution for implementing network storage automation. These systems can be segmented into three general categories: internal stackers/ loaders, desktop libraries and large floor-standing libraries. Desktop libraries using high capacity 4 mm and 8 mm tape systems typically support capacities of 100 Gbytes or more. Although floor-standing libraries are the obvious choice for very large networks with storage requirements of several hundred Gbytes, desktop libraries address the widest range of general network storage needs.

Stackers vs Libraries

As with most emerging markets, automated network storage can be a confusing domain of new terminology and blurred product differentiation. Some vendors of low-end stacker and loader mechanisms have attempted to position their products as true libraries. However, there are two clear distinctions that separate stackers from libraries. First, a library will accommodate a significantly higher number of media than stackers. Stackers or loaders are commonly built with an internal drive as part of a 5.25-in. form factor device, limiting the capacity to four to six tapes. Desktop libraries offer at least twice the media capacity, while larger floor models can contain hundreds of cartridges.

Second, libraries are distinguished by their ability to randomly access any tape cartridge with complete software control of the process. The most common library design provides separate SCSI addresses for the drive and the robotics mechanism, providing maximum flexibility in configuring the library for use on a SCSI bus. Stackers are constructed with only a single SCSI address so that any command going to the mechanism must be passed through by the drive, making it impossible for the drive and mechanism to receive data simultaneously.

True libraries are also much more reliable designs than stackers/loaders. Today's desktop libraries include robotic changer mechanisms that will likely outlast the tape drives included in the unit. The minimum acceptable standard for the changer is nominally 100,000 media swaps between failure, particularly in random search operations. Stacker designs don't approach that level of reliability, and likely never will as these devices were never intended for intense random access operations and as a result use less robust mechanics in the changer mechanisms.

The complexity of stacker/ loader designs also makes these changers more expensive to repair, both in terms of the physical repair and the down time involved. The changer mechanism and the tape drive are a single device, making it more difficult and time consuming to repair. A jammed mechanism in an internal stacker could result in having to bring the file server down, disrupting the network. Library designs allow a simple drive swap that takes 30 minutes or less. In dual-drive libraries, the repair can be done on the first drive while the second drive continues to operate so that the library is never out of service.

The more innovative library architectures include dual tape drives that allow the first drive to write information coming from the server while the robotic changer is instructed to retrieve a tape so that a file can be read back at the same time on the second drive. A dual-drive library also supports direct tape-to-tape copies, a key feature in producing tapes for secure offsite archival storage. This level of functionality is impossible with stackers and loaders unless a second unit is purchased.

Library Advantages

The attraction of automated libraries is easy to see. Adding a desktop library to a network storage system offers three principal benefits: improved data security, increased organization of storage media, and significant cost savings.

Perhaps the most attractive feature of a desktop tape library is that, coupled with appropriate software, it eliminates human errors from the data security equation. Repeated studies have demonstrated that human error, not hardware failure, is the common cause of data loss. Programming a software application to automatically backup or migrate files off the hard disk to tape based on specified criteria ensures that the routine will be performed perfectly every time, protecting critical information.

The switch to automated storage libraries can virtually eliminate media storage and organization problems. With dozens of cartridges loaded into the library, network managers no longer have to go searching through several individual tapes trying to locate a single file that was removed from the server but has been requested by a network user. Virtually all critical data are now available on the near-line library. Under HSM software control the automated storage system will prompt the administrator for the specific tape to load if it is not found among the current library files.

The best library designs further add to organizational simplicity by keeping a full data set of multiple tapes together in a single magazine. Instead of the administrator loading individual tapes into slots in the library, the media magazine is pre-loaded with tapes prior to being inserted into the library. Today's backup and migration sessions can easily span several tapes. Using the magazine approach, all of the tapes relating to a specific backup session are stored together, eliminating any chance of misplacing any individual tape. Media magazines also provide a high level of data protection. With a secure dust cover in place, magazines are the ideal transportation medium when moving tapes to offsite vault storage. The magazine provides a contaminant-free environment both during transportation and shelf storage of critical data.

The most persuasive argument in favor of adding an automated library to a network storage environment is clearly cost savings. A 4 mm-based automated desktop library can be purchased for around $7,500 with a hardware cost of just 11 cents/Mbyte. Based on the above-mentioned estimate of network administration—as high as $40,000 a year for 2 Gbytes of data—the system could pay for itself in months. Moreover, the automation capability requires very little additional training for the administrator and is based on tape technologies that are well-proven and with which LAN managers are comfortable.■

Dan Stromska *is vice president of sales and marketing for ADIC (Redmond, WA).* **For more info circle 107.**

Advertisers' Index

	Page	Reader Inquiry Card
3M Data Storage Products	97	28
Ancot Corporation	96	27
Anritsu	15	63
CMD Technology, Inc.	67	18
Colorado Memory Systems	Cover 4	2
Computone Corp.	22	7
Concurrent Computer Corp.	12	6
Connector Resources Unlimited	58	66
Data General Corp.	2-3	3
Deltec	31	38
Distributed Processing Technology, Inc.	Cover 3	1
Epson America, Inc.	60-61	16
Exabyte Corp.	102-103	32
Gigatek Memory Systems	90	24
Hitachi America Ltd.	1	43
HT Research	4	4
IBM	Cover 2	55
Imperial Technology, Inc	87	21
IT Times	35	62
JMR Electronics, Inc.	37	39
JMR Electronics, Inc.	69-80	–
M4 Data, Inc.	100	30
Merisel	11	65
Merisel	20	60
Metrum	92-93	10-15
Micro Design International, Inc.	53	80
Micro House International	43	40
Olympus Image Systems, Inc.	55	42
Optima	25	61
Overland Data Inc.	95	26
Pentax	47	46
Philips Key Modules Group	45	45
Philips-LMS	110	36-37
Pioneer New Media Technologies, Inc.	41	44
Rexon	99	29
Rexon	107	34
Rexon	109	35
Sankyo Seiki	104	33
SCM Microsystems, Inc.	85	64
Seagate Technology	9	5
Sigma Information Systems	62	17
Sony Recording Media	101	31
Star Gate Technologies, Inc.	27	8
Storage Computer Corp.	88	22
Technology Distributors, Inc.	89	23
Trimm Industries	82	19
Upsonic	28	9
Verbatim Corp.	50	41

This advertisers' index is an additional service provided by CTR. The Publisher assumes no liability for omissions or errors.

COMPUTER TECHNOLOGY REVIEW®

READER INQUIRY CARD Receive free information on the products featured in this issue. Just circle the reader inquiry numbers below.

LABEL ID # (5 OR 6 DIGITS ABOVE NAME) _____

NAME _____ TITLE _____

COMPANY _____

ADDRESS _____

CITY _____ STATE _____

TELEPHONE _____ ZIP _____

COMMENTS: _____

If this is not your personal subscription:
☐ check here to receive qualification information.

01	16	31	46	61	76	91	106	121	136	151	166	181	196	211	226	241	256	271	286
02	17	32	47	62	77	92	107	122	137	152	167	182	197	212	227	242	257	272	287
03	18	33	48	63	78	93	108	123	138	153	168	183	198	213	228	243	258	273	288
04	19	34	49	64	79	94	109	124	139	154	169	184	199	214	229	244	259	274	289
05	20	35	50	65	80	95	110	125	140	155	170	185	200	215	230	245	260	275	290
06	21	36	51	66	81	96	111	126	141	156	171	186	201	216	231	246	261	276	291
07	22	37	52	67	82	97	112	127	142	157	172	187	202	217	232	247	262	277	292
08	23	38	53	68	83	98	113	128	143	158	173	188	203	218	233	248	263	278	293
09	24	39	54	69	84	99	114	129	144	159	174	189	204	219	234	249	264	279	294
10	25	40	55	70	85	100	115	130	145	160	175	190	205	220	235	250	265	280	295
11	26	41	56	71	86	101	116	131	146	161	176	191	206	221	236	251	266	281	296
12	27	42	57	72	87	102	117	132	147	162	177	192	207	222	237	252	267	282	297
13	28	43	58	73	88	103	118	133	148	163	178	193	208	223	238	253	268	283	298
14	29	44	59	74	89	104	119	134	149	164	179	194	209	224	239	254	269	284	299
15	30	45	60	75	90	105	120	135	150	165	180	195	210	225	240	255	270	285	300

READER SERVICE NUMBERS APPLY TO ITEMS APPEARING WITHIN THIS ISSUE ONLY!

BS4C

COMPUTER TECHNOLOGY REVIEW®

READER INQUIRY CARD Receive free information on the products featured in this issue. Just circle the reader inquiry numbers below.

LABEL ID # (5 OR 6 DIGITS ABOVE NAME) _____

NAME _____ TITLE _____

COMPANY _____

ADDRESS _____

CITY _____ STATE _____

TELEPHONE _____ ZIP _____

COMMENTS: _____

If this is not your personal subscription:
☐ check here to receive qualification information.

01	16	31	46	61	76	91	106	121	136	151	166	181	196	211	226	241	256	271	286
02	17	32	47	62	77	92	107	122	137	152	167	182	197	212	227	242	257	272	287
03	18	33	48	63	78	93	108	123	138	153	168	183	198	213	228	243	258	273	288
04	19	34	49	64	79	94	109	124	139	154	169	184	199	214	229	244	259	274	289
05	20	35	50	65	80	95	110	125	140	155	170	185	200	215	230	245	260	275	290
06	21	36	51	66	81	96	111	126	141	156	171	186	201	216	231	246	261	276	291
07	22	37	52	67	82	97	112	127	142	157	172	187	202	217	232	247	262	277	292
08	23	38	53	68	83	98	113	128	143	158	173	188	203	218	233	248	263	278	293
09	24	39	54	69	84	99	114	129	144	159	174	189	204	219	234	249	264	279	294
10	25	40	55	70	85	100	115	130	145	160	175	190	205	220	235	250	265	280	295
11	26	41	56	71	86	101	116	131	146	161	176	191	206	221	236	251	266	281	296
12	27	42	57	72	87	102	117	132	147	162	177	192	207	222	237	252	267	282	297
13	28	43	58	73	88	103	118	133	148	163	178	193	208	223	238	253	268	283	298
14	29	44	59	74	89	104	119	134	149	164	179	194	209	224	239	254	269	284	299
15	30	45	60	75	90	105	120	135	150	165	180	195	210	225	240	255	270	285	300

READER SERVICE NUMBERS APPLY TO ITEMS APPEARING WITHIN THIS ISSUE ONLY!

BS4C

All payments in U.S. funds only. You may fax your Visa order: 310-208-1054

COMPUTER TECHNOLOGY REVIEW®

☑ YES! I need the technology publication for the integration/reseller market. Send me Computer Technology Review today! 16 data-packed issues for only: $80 US/ $100 Canada/ $160 All Others. If I'm not happy with CTR I can cancel at anytime and receive a full refund for the unmailed issues.

❑ My check or money order is enclosed.

❑ Please charge my VISA card # _____ Expire Date _____

Signature _____

NAME _____ **TITLE** _____

COMPANY _____

MAILING ADDRESS _____

_____

_____ **COUNTRY** _____

Name _____
Company _____
Address _____
City/State/Zip _____

PLEASE PLACE POSTAGE HERE

COMPUTER TECHNOLOGY REVIEW

924 Westwood Blvd., Suite 650
Los Angeles, CA 90024-2910

Name _____
Company _____
Address _____
City/State/Zip _____

PLEASE PLACE POSTAGE HERE

COMPUTER TECHNOLOGY REVIEW

924 Westwood Blvd., Suite 650
Los Angeles, CA 90024-2910

HERE'S HOW:

- ✓ Complete form on reverse side.
- ✓ Enclose check or money order (U.S. funds only) OR charge to your VISA card.
- ✓ Mail today! Subscriptions begin with current issue upon payment, or specify issue within current year.

DON'T MISS THESE UPCOMING FEATURES:

- Full Motion Video Over Networks
- Wireless LANs
- Windows NT or UNIX
- Hierarchical Storage Management
- Multimedia
- Distributed Processing
- PCMCIA Developments
- X Stations
- and much, much more!

SUBSCRIBE TODAY!

COMPUTER TECHNOLOGY REVIEW

924 Westwood Blvd., Suite 650
Los Angeles, CA 90024-2910 USA

Card will fit in #10 window envelope.

1. *Which best describes your company?* (select one)

SYSTEM INTEGRATOR/THIRD PARTY RESELLER/VAR

A. Reseller of Computer Systems/Networks/VAR
B. System Integrator/Systems House/VAR
C. Software Integrator/Software House/VAR
D. Distributor/Dealer/Master VAR
E. Consultant
F. Government (Federal, State, Local)
G. Military/Aerospace
H. Other System Integrator

INSERT CODE

OEM/SYSTEMS INTEGRATOR

I. Manufacturer/Integrator of Computer/Computer Systems
J. Manufacturer/Integrator of Peripherals
K. Manufacturer/Integrator of Networking Products/Comm Equipment
L. Other OEM/Mfg/Industrial Which Integrates/Resells

IN-HOUSE SYSTEMS INTEGRATOR IN:

M. Manufacturing/Industrial/Commercial/Utilities
N. Financial/Legal/Health Services/Insurance
O. Government (Federal, State, Local)
P. Military/Aerospace
Q Retail/Wholesale/Distribution
R. Education/Scientific
S. Other in-house integrator

None of The Above (specify) _____

2. *Which best describes your job function?* (select only one)

CORPORATE MANAGEMENT

A. Executive Management (Chairman, President, Partner, etc.)
B. Corporate Management (VP, General Manager, Director, etc.)

TECHNICAL/ENGINEERING MANAGEMENT/STAFF

C. Engineering Management
D. Systems Integrator/Consultants
E. Systems Consultants/Systems Analysts
F. Software Management/Staff
G. Network Management/Staff
H. Information Systems Management/Staff
I. Systems Design/Consultant
J. Software Development/Consultant
K. Technical Management/Staff
L. Marketing/Sales Staff

INSERT CODE

None of The Above (specify) _____

3. *What is the annual business volume of your company?* (select only one)

A. $1 Billion and Over
B. $100 Million to $999.9 Million
C. $50 Million to $99.9 Million
D. $25 Million to $49.9 Million
E. $10 Million to $24.9 Million
F. $1 Million to $9.9 Million
G. $500,000 to $999,999
H. Under $500,000

INSERT CODE

4. *What is the approximate annual dollar volume of purchases you company will make and/or specify for systems/networking integration?* (select only one)

A. Over $50 Million
B. $25 Million to $49.9 Million
C. $10 Million to $24.9 Million
D. $1 Million to $9,999,999
E. $500,000 to $999,999
F. Under $500,000

INSERT CODE

5. *What operating systems does your company integrate, specify, or resell systems based on?* (select only one)

OPERATING SYSTEMS	LAN OPERATING
A. ❑ UNIX/UNIX Flavors	H. ❑ Banyan (Vines)
B. ❑ DOS	I. ❑ Microsoft (LAN Manager)
C. ❑ OS/2	J. ❑ Novell (Netware)
D. ❑ Apple/MacIntosh	K. ❑ Apple Talk/Share
E. ❑ VAX/VMS	L. ❑ Artisoft (LANtasic)
F. ❑ Windows	M. ❑ Windows for Workgroup
FF. ❑ Windows NT	Other (Specify)_____
G. ❑ Multiuser/DOS	
GG. ❑ DESQview/DVX	

Other (Specify)_____

6. *What type of your environment do you automate?* (select only one)

BY SYSTEM SIZE

A. ❑ Large scale systems (more than 128 users)
B. ❑ Medium scale systems (17-128 users)
C. ❑ Small scale systems (1-16 users)
Z. ❑ Not Applicable

BY SIZE OF COMPANY SOLD TO

D. ❑ Large (Over 1,000 employees)
E. ❑ Medium (100-1,000 employees)
F. ❑ Small (Under 100 employees)
Z. ❑ Not Applicable

7. *You resell, purchase, specify, recommend or approve the purchase of, which of the following products:* (select only one)

COMPUTERS/SYSTEMS	MASS STORAGE
A. ❑ Mainframe/Supercomputers	A. ❑ Hard Disk Drives
B. ❑ Minicomputers/Superminis	B. ❑ Flexible Disk Drives
C. ❑ Personal/Desktop Computers	C. ❑ Solid State Disks/Flash
D. ❑ Laptop/Notebook/Pen Computers	D. ❑ Disk Arrays (RAID)
E. ❑ Multimedia Systems	E. ❑ Disk Subsystems
F. ❑ Embedded Systems	F. ❑ Erasable Optical Drives
G. ❑ Workstations	G. ❑ WORM Drives
PERIPHERALS	H. ❑ CD-ROM Drives
H. ❑ Bus Boards (VME/EIS etc.)	I. ❑ Optical Subsystems
I. ❑ Interface Boards (SCSI/ESDI etc.)	J. ❑ DAT-4mm Tape Drives
IX. ❑ PCMCIA Interface	K. ❑ 8mm Tape Drives
J. ❑ Multimedia Boards	L. ❑ 1/4" Tape Drives
K. ❑ Controllers	M. ❑ 1/2" Tape Drives
L. ❑ Monitors/Displays	N. ❑ Tape Subsystems
M. ❑ XTerminals/Diskless Workstations	NX. ❑ Media Optical/Tape
N. ❑ Printers	**SOFTWARE**
O. ❑ Scanners/Input Devices	O. ❑ Real Time
P. ❑ UPS/SPS	P. ❑ Networking Software
NETWORKING/COMMUNICATIONS	Q. ❑ RDBMS
Q. ❑ LANs	R. ❑ ODBMS
R. ❑ Multiuser Systems	S. ❑ Multimedia
S. ❑ Servers	T. ❑ GUIs
T. ❑ Internetworking Equipment (Bridges, Routers, Gateways)	V. ❑ Graphics Packages
	W. ❑ Expert/AI Packages
V. ❑ Network Power Security Equip.	Y. ❑ HSM
W. ❑ Wireless WANs	Z. ❑ Document Imaging Mgmt.

❑ **None of the above**

8. *How do you purchase your hardware and software products?* (select only one)

	HARDWARE	SOFTWARE
Direct from Manufacturer	A. ❑	E. ❑
From Distributor	B. ❑	F. ❑
From Super/Master VAR	C. ❑	G. ❑
From other 3rd Party	D. ❑	H. ❑

Other (specify)

9. *Which of the following vendors' products does your company integrate, use, specify, or resell?* (select all that apply)

SOFTWARE	HARDWARE
A. ❑ Autodesk	A. ❑ ACER/ALTOS
B. ❑ Banyan	B. ❑ ALR
C. ❑ Borland	C. ❑ Apple/MacIntosh
D. ❑ Claris	D. ❑ AST
E. ❑ Cognos	E. ❑ AT&T/NCR
F. ❑ Concurrent Controls	G. ❑ Compaq
G. ❑ Digitalk	H. ❑ Data General
H. ❑ Gupta	J. ❑ DEC VAX/Workstations
I. ❑ Informix	K. ❑ Dell
J. ❑ Interactive	L. ❑ Epson
K. ❑ Lotus	M. ❑ Everex
L. ❑ Microsoft	O. ❑ HP
M. ❑ Novell	P. ❑ IBM Workstations
N. ❑ Oracle	Q. ❑ Other IBM Compatibles
O. ❑ Parc Place	R. ❑ Intergraph
P. ❑ Performance Technology	S. ❑ Leading Edge
Q. ❑ Progress	T. ❑ Motorola
R. ❑ Ready Systems	U. ❑ NEC
S. ❑ SCO	W. ❑ Samsung
T. ❑ Servio	X. ❑ Silicon Graphics
U. ❑ Software Link	XY. ❑ Sony
V. ❑ Sybase	Y. ❑ Sun
W. ❑ Theos	Z. ❑ Texas Instruments
X. ❑ Unify	AA. ❑ Toshiba
Y. ❑ Uniplex	BB. ❑ Unisys
Z. ❑ Versant	CC. ❑ Wyse
AA. ❑ WindRiver	DD. ❑ Zenith Data Systems
BB. ❑ Word Perfect	
CC. ❑ Quarterdeck	
Other (specify)_____	Other (specify)_____

❑ **None of the above**

10. *Which market/application areas does your firm integrate/resell products into?* (select only one)

A. ❑ Business Services/Office Automation/DP
B. ❑ Financial/Insurance/Health Organizations/Legal
C. ❑ Manufacturing/Industrial/Utilities
D. ❑ Education
E. ❑ Government (Federal,State,Local)
F. ❑ Aerospace/Military
G. ❑ Scientific/Engineering (including CAD/CAE)
H. ❑ Graphics/Imaging/Multimedia
I. ❑ Networking
J. ❑ Retail/Wholesale/Distribution
L. ❑ Fortune 1000
Other (specify)_____

11. *In the purchasing process, your direct responsibilities are to:* (select all that apply)

A. ❑ Evaluate the technology
B. ❑ Review/Select the brand/Manufacturer
C. ❑ Approve the purchase
Z. ❑ No direct responsibility

12. *In addition to COMPUTER TECHNOLOGY REVIEW, what three industry publications do you read?*

_____

_____

Turn this page over and answer all questions to start your free subscription

COMPUTER TECHNOLOGY REVIEW®

FREE SUBSCRIPTION APPLICATION

Please sign and date your subscription request. Incomplete forms cannot be processed. The Publisher reserves the right to accept only those applications which met the qualification criteria for the publication.

❑ **YES!** I want to receive COMPUTER TECHNOLOGY REVIEW

❑ No

Signature_____ **Date**_____
Required

Phone(____) _____ **Fax**(____) _____

❑ ***New Subscription***
Print legibly

Name _____

Title _____

Company _____

Street Address _____

City _____ **State**____ **Zip** _____

If you would prefer delivery to your home, please complete below. Company name and address still required above.

Home Address _____

City _____ **State**____ **Zip** _____

(Free subscriptions are within the U.S.A. and territories)
Canada, Mexico $100 All other countries $160, Airmail

PS4C

Fold here

PLEASE PLACE POSTAGE HERE

COMPUTER TECHNOLOGY REVIEW
924 Westwood Blvd., Suite 650
Los Angeles, CA 90024-2910

Fold here

MAIL CARD AS IS.
NO ENVELOPE REQUIRED.

TAPE HERE ONLY

MCI Relies On DPT SCSI Adapters...

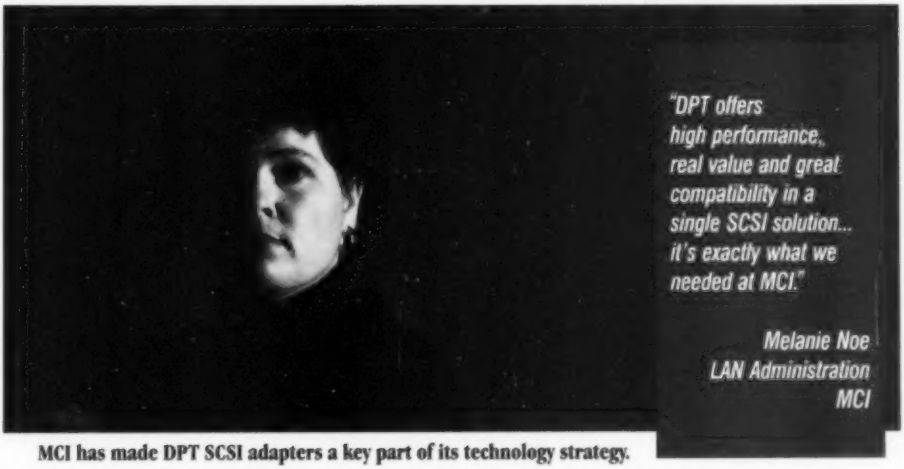
"DPT offers high performance, real value and great compatibility in a single SCSI solution... it's exactly what we needed at MCI."
Melanie Noe
LAN Administration
MCI

MCI has made DPT SCSI adapters a key part of its technology strategy.

As one of today's leading technology companies, MCI knows that high-performance servers and workstations need high disk throughput, compatibility, and flexibility for growth. That's why MCI relies on DPT SCSI adapters as an integral part of its networked systems.

"Compatibility, reliability, and performance were key factors in our decision process in choosing a SCSI host adapter. DPT outperformed everything else we tested in these critical categories." That's because DPT adapter support is integrated into every major multi-user Operating System available today, including Novell, UNIX, OS/2, Banyan VINES, NEXTSTEP, and Windows NT. Full support for a complete range of important SCSI peripheral products and software is also included.

"At MCI, the only certainty is that things are certainly going to change! That's why we like DPT's growth path. With their unique approach, we can add extra capabilities such as caching or RAID support, whenever we want." Attach a cache module and add up to 64MB of hardware cache using industry-standard or DPT's ECC-protected SIMMs. Then add RAID 0, 1, or 5 support by simply attaching another easy-to-install module. It's a proven approach designed to keep pace with the growth and changes in your system, and only DPT has it.

DPT's new **Storage Manager**, called *"a new standard in storage management"* by LAN TIMES, is included **free** with every DPT SmartCache III adapter. Storage Manager gives you complete device inventory and configuration data, diagnostic and performance tuning abilities, capabilities for remote communications and alert broadcasting, along with the ability to implement complex RAID structures with simple "point-and-click" commands.

DPT supplies SCSI technology to the world's largest computer manufacturers, and has an unsurpassed record for product reliability among the foremost distributors in over 40 countries. Call today and receive a free copy of "Understanding RAID" and "Hardware and Software Caching Combined."

DPT
for complete SCSI solutions

800-322-4378

Distributed Processing Technology · 140 Candace Drive · Maitland, Florida · 32751 · FAX 407-260-5366

yes NetWare Tested and Approved · BANYAN · SCO OK

Circle 1 on Reader Inquiry Card

COLORADO TAPE BACKUP

EXPLORE COLORADO
PowerTape Series 1100
1 GB
HIGH COUNTRY

4 Gigabytes.
Under $1,700.
POWERTAPE

When you consider the value of all the data on a network, any backup system is a wise investment. But now you have a choice. PowerTape Series 4000 from Colorado Memory Systems. With a native capacity of 2 GB and 4 GB using data compression, PowerTape can handle virtually any backup task.

The PowerTape 4000 system includes a backup tape drive, SCSI controller, Colorado Backup™ software, data cartridge, and more. At $1,695 MSRP why pay more and get less? And, if 4 GB is more than you need, a 2.4 GB model is available with the same high peformance features at only $1,295 MSRP. And a new low-cost 1 GB model is available. At only $799 MSRP, it's the most affordable PowerTape ever.

With PowerTape, you get easy restores and backups, plus the reliability of QIC Industry Standard Format. PowerTape offers broad network compatibility including Novell® NetWare®, NetWare Lite™, LANtastic®, and IBM® PC-NET and operates under UNIX®, OS/2® and Microsoft® Windows NT™ operating environments. And all PowerTape systems are backed by the best support in the industry.

Call **1-800-451-0897** extension 3111 today for a FREE Colorado *Total Backup* catalog.

POWERTAPE™
SERIES 4000

COLORADO MEMORY SYSTEMS INC.

* All native capacity. © 1994 Colorado Memory Systems, Inc. All rights reserved.
Colorado and Jumbo are trademarks of Colorado Memory Systems Inc. All other trademarks belong to their respective companies. 94MCTB11295

Circle 2 on Reader Inquiry Card